

Behavioral Mechanism Design: Optimal Contests for Simple Agents

Arpita Ghosh and Robert Kleinberg
Cornell University
Ithaca, NY, USA

Abstract

Incentives are more likely to elicit desired outcomes when they are designed based on accurate models of agents' strategic behavior. A growing literature, however, suggests that people do not quite behave like standard economic agents in a variety of environments, both online and offline. What consequences might such differences have for the optimal *design* of mechanisms in these environments? In this paper, we explore this question in the context of optimal contest design for *simple* agents—agents who strategically reason about whether or not to participate in a system, but not about the input they provide to it. Specifically, consider a contest where n potential contestants with types (q_i, c_i) each choose between participating and producing a submission of quality q_i at cost c_i , versus not participating at all, to maximize their utilities. How should a principal distribute a total prize V amongst the n ranks to maximize some increasing function of the qualities of elicited submissions in a contest with such simple agents?

We first solve the optimal contest design problem for settings where agents have homogenous participation costs $c_i = c$. Here, the contest that maximizes every increasing function of the elicited contributions is always a *simple* contest, awarding equal prizes of V/j^* each to the top $j^* = V/c - \Theta(\sqrt{V/(c \ln(V/c))})$ contestants. This is in contrast with the optimal contest structure in comparable models with strategic effort choices, where the optimal contest is either a winner-take-all contest or awards possibly unequal prizes, depending on the curvature of agents' effort cost functions. We next address the general case with heterogenous costs where agents' types (q_i, c_i) are inherently two-dimensional, significantly complicating equilibrium analysis. With heterogenous costs, the optimal contest depends on the objective being maximized: our main result here is that the winner-take-all contest is a 3-approximation of the optimal contest when the principal's objective is to maximize the quality of the best elicited contribution. The proof of this result hinges around a 'sub-equilibrium' lemma establishing a stochastic dominance relation between the distribution of qualities elicited in an equilibrium and a *sub-equilibrium*—a strategy profile that is a best response for all agents who choose to *participate* in that strategy profile; this relation between equilibria and sub-equilibria may be of more general interest.

1 Introduction

A vast range of online systems—social computing and crowdsourcing platforms, matching markets, ranking and recommendation systems—involve agents who make choices that determine the input to, and therefore the outcome of, the system. This poses an immediate problem of designing incentives and environments that drive choices resulting in desirable outcomes, and has resulted in a rich literature on mechanism design and algorithmic game theory for these settings.

Economic design is more likely to be effective when it is based on accurate models of the agent population at which it is targeted. A growing literature, however, suggests that people do not quite behave like

standard economic agents in a variety of settings, both online and offline. What consequences might such differences in behavior have for the optimal *design* of these environments?

In this paper, we explore this question of ‘behavioral’ mechanism design in the context of designing optimal contests for *simple* agents—agents who strategically reason about whether or not to participate in a system, but not about the input they provide to it—to see how this difference in agent behavior affects the analysis of the strategic environment and the structure of optimal mechanisms.

Before proceeding, we note that our use of the term ‘behavioral’ differs slightly from that in the behavioral economics literature, and specifically is not meant to *require* that our agents are irrational in any sense or have bounded reasoning abilities: indeed, the simple agents in our model do not make suboptimal choices—they do make optimal choices, but over restricted *choice* sets relative to the standard models in the contest design literature. We use the term behavioral only to indicate (mechanism design problems with) models of agents that more accurately capture real human behavior in the environment of interest than standard economic models.

1.1 Contest design

Contests, where participants expend some resource—time, effort, money—to compete for prizes, are everywhere. There are athletic and artistic competitions and competitions for mathematics and research development and design; online, too, there are an exploding number of contests—both explicit, like online photography contests and design contests on platforms like Quirky and 99designs¹, as well as implicit contests for attention or virtual rewards² on user-contribution based websites. The ubiquitousness of contests, as well as the seminal work of Lazear and Rosen [19] showing that contests can serve as an efficient incentive scheme for effort elicitation, has led to a huge literature on contest design, studying how to allocate rewards to best incentivize high-quality submissions from utility-maximizing contestants with a cost to effort.

There is a wide spectrum of contests, however, where participants do not quite strategize over the quality of their submissions:

- First, there are contests where the *production* of an entry is essentially costless, and the only cost incurred is that of submission—for instance, potential contestants may already possess eligible entries prior to, or independent of, their knowledge of the contest, and decide only whether or not to incur the cost (either effort or monetary) of submission, as in arts or literature contests³.
- Second, there are environments where producing an entry indeed incurs a cost, but this cost essentially does not vary with the quality of the entry—for example, contributing an answer on expertise-based⁴ online Q&A forums incurs a cost⁵ which is essentially independent of the contribution’s quality.
- Finally, there are ‘best-effort’ environments where producing an entry does incur a cost, and incurring a higher effort cost does increase the quality of the resulting submission, but where contestants do

¹www.quirky.com; www.99designs.com

²Most user-generated content sites rank users’ contributions according to some measure of quality, and a contribution’s ranking determines its likely viewership and therefore the attention reward to the corresponding user. A number of websites also give out more explicit virtual rewards to encourage contribution, such as virtual points (e.g. Y! Answers), badges (StackOverflow, Quora, TripAdvisor) or leaderboard rankings (Amazon, Y! Answers).

³As one example, a number of travel photography contests are frequently hosted by both government agencies and private enterprises such as National Geographic—participants in these contests likely do not undertake their travel to produce a photo for such a contest, but rather choose whether to submit an eligible photo or not.

⁴(i.e., where the quality or value of an answer depends much more on the expertise of the answerer than her effort, such as health or law forums)

⁵(the cost of logging in and posting the answer on the forum)

not ‘game’ their effort—having made the decision to participate, a contestant will work to the best of her ability to produce the best entry she possibly can. Any contest-like scenario where entry is driven by intrinsic motivation for the task at hand can fall into this category. This category includes offline contests for, say, a performing art. It also includes online contests like Quirky, as well as citizen-science and user-contributed content driven communities. Indeed, there is a growing literature suggesting that users in several online contribution settings do not strategize on effort on a task in response to incentives, although they might choose whether or not to undertake the task based on the promised rewards [17]⁶; for instance, [16] find in a field experiment on Google Answers that participation, but not quality, is sensitive to the prize offered to winning answers.

There are clearly a vast range of real-life scenarios, therefore, where potential contestants do strategize about whether or not to participate in a contest, but not about the quality of their submissions. How do incentives and equilibrium outcomes change when agents only make strategic participation choices, rather than strategic effort choices, to maximize their expected utilities? Specifically, what allocation of prizes for each rank leads to the ‘best’ elicited set of submissions in equilibrium?

1.2 Our contributions

We explore the idea of ‘behavioral’ mechanism design via the problem of designing optimal contests for *simple* agents, where a principal seeks to allocate a total prize to optimize the qualities of elicited submissions when potential contestants only strategize about whether to participate in the contest. Specifically, consider a population of n potential contestants with types (q_i, c_i) drawn from a known joint distribution $F(q, c)$, where an agent with type (q_i, c_i) produces output of quality q_i at a cost c_i if she participates in the contest. Suppose there is a principal who can distribute a total prize V amongst agents based on their relative ranks; here V might be either a monetary prize budget, or a non-monetary resource such as display space on a webpage that translates to attention rewards. Each agent decides whether or not to participate by comparing her expected prize from participating (which depends on the rank of her submission’s quality q_i relative to that of the other agents who also choose to enter) against her cost of participation c_i . What allocation of V amongst the n ranks induces equilibrium participation decisions that lead to the best outcomes?

Participation-only strategic choices can be thought of as a special case of a strategic effort choice model, where agents are restricted to a binary choice between the maximum possible effort and zero effort. While it might seem, at first glance, that reducing agents’ choices from an interval to the endpoints of that interval ought not change the nature of the optimal contest design problem much, or at least only make it simpler, this turns out to be far from the truth—agents’ incentives are significantly altered by the fact that any competitor who does choose to participate will produce her ‘best possible output’ instead of ‘adjusting’ her effort (and correspondingly the submission quality that all other participants must compete against) to some suitable point in the interior of the interval. We note here an interesting conceptual parallel with discrete versus continuous optimization: moving from continuous to discrete optimization problems where variables can only take binary (or integer) values fundamentally changes the nature of the algorithm design question, leading to an entirely different body of techniques for algorithm design and, in many cases, to computationally intractable problems. A second factor that contributes to the complexity of the strategic problem in these settings is that agents can, in general, have two-dimensional types: as discussed below, this

⁶Kraut and Resnick [17] summarize this as a ‘design claim’: “With task-contingent rewards for small, discrete tasks, larger rewards will motivate people to take on tasks, but will not motivate higher effort on accepted tasks.”

significantly complicates the nature of equilibria, necessitating new tools to bound the outcomes elicited in an equilibrium of a contest.

We first study the optimal contest design problem in settings with homogenous participation costs c , so that agents' types (q_i, c) are essentially one-dimensional: here, the contest that elicits the highest equilibrium participation is optimal for any objective that is an increasing set function of the elicited submission qualities. We show that this optimal contest is always a *simple contest* which awards equal prizes of V/j^* each to the top j^* contestants, where the optimal number of prizes j^* scales with $\nu = V/c$ as $\nu - \Theta(\sqrt{\nu/\log(\nu)})$. This contrasts with the optimal contest structure in comparable models with strategic effort choices, where the optimal contest is either a winner-take-all contest or may award unequal prizes, depending on the curvature of agents' effort cost functions [21]. This means, for example, that in online user-generated content environments where users do decide whether or not to contribute, but do not strategize over the qualities of the content they produce, handing out an appropriately chosen number of identical badges might lead to better outcomes than reward structures—such as ranked top-contributor lists—inducing unequal social-psychological rewards. As another example, a suitably chosen number of equal 'merit' prizes in a travel photography contest might elicit superior submissions than a structure that awards unequal first, second, and third prizes⁷.

We next address the general case when agents have heterogenous costs c_i that can be arbitrarily correlated with the qualities q_i of their submissions. Here, agents' types (q_i, c_i) are inherently two-dimensional, unlike in the homogenous cost model in §3, or in strategic effort models in the contest design literature (where agents' types are parametrized by a one-dimensional ability that either scales agents' cost functions or their outputs per unit effort). With two-dimensional types, much of the structure that is central to the equilibrium analysis of contests with one-dimensional agent types—where equilibrium action choices are typically monotone in this type—vanishes, making it difficult to explicitly characterize equilibria, and therefore to identify the contest structure that yields the 'best' equilibrium outcome.

Since the design of the optimal contest can, in general, vary depending on the choice of objective function in this setting, we focus on contests that maximize the quality of the *best* elicited submission in equilibrium. Our main result in this model concerns the winner-take-all contest, which awards the entire prize V to the top entry. A winner-take-all contest does not necessarily optimize the expected quality of the best elicited submission, as shown, for example, by our analysis of the case of homogeneous participation costs: we show, however, that the winner-take-all contest always achieves a bounded approximation factor⁸ with respect to the optimal contest. More precisely, the expected value of the highest-quality entry elicited in any pure-strategy equilibrium of the winner-take-all contest is no smaller than one-third of the expected maximum quality elicited in the best equilibrium of an optimal contest \mathcal{M}^* . One interpretation of this result is that the designer's problem has an approximately detail-free solution: without knowing the joint distribution of (q_i, c_i) , the designer can run a winner-take-all contest and thereby guarantee, in equilibrium, an outcome not much worse than the best equilibrium achievable by any contest. The proof of this result hinges around

⁷For example, the National Geographic 2014 contest awards 3 unequal top prizes and 7 identical merit prizes: <http://travel.nationalgeographic.com/travel/traveler-magazine/photo-contest/2014/>.

⁸One could question whether such approximation results are meaningful, since it is unclear what it even means to quantify the quality of elicited submissions numerically. Our approximation result is at least partly immune from this critique for the following reason. First note that the critique applies only to the way we evaluate the quality of a mechanism's outcome in equilibrium, not to the way we model mechanisms, agents, and equilibria. This is because the mechanisms we consider are rank-order contests which make only ordinal comparisons of quality, so the behavior of a mechanism and its set of equilibria are unaffected by monotonic reparameterizations of the quality scale. For essentially the same reason, our approximation result about the winner-take-all contest remains valid under any monotonic reparameterization of quality; the reparameterization does not affect the set of equilibria of the contest, so the conclusion of any theorem about that set of equilibria cannot be sensitive to whether we have reparameterized the quality scale.

a ‘sub-equilibrium’ lemma, which establishes a stochastic dominance relation between the distribution of qualities elicited in an equilibrium and in a *sub-equilibrium*—a strategy profile that is a best response for all agents who choose to *participate*, but not necessarily a best response for non-participants—of the winner-take-all contest, and may have broader applications to characterizing equilibrium performance in settings where it is difficult to explicitly solve for equilibria.

1.3 Related work

There is now a tremendously large literature on the economics of contests. The seminal work of Lazear and Rosen [19] suggesting that contests, or rank-order tournaments, can serve as an efficient scheme to incentivize effort from strategic agents has given rise to a vast literature on contests as a means for effort elicitation, see e.g. [13, 5]. The problem of optimal contest design—how to allocate rewards to elicit the most desirable outcomes, posed by Sir Francis Galton back in 1902 and formally addressed as early as [11]—has been studied in a vast variety of settings, encompassing heterogeneous (e.g. [21, 22, 4]) and homogenous (e.g. [26, 10]) agent populations, risk-neutral and risk-averse preferences over prizes, non-monetary rewards [23, 10], as well as various models of information and observability of output, and various objectives and constraints faced by the principal running the contest.

While very large, this literature—to the best of our knowledge—studies contests where agents make strategic effort choices that (deterministically or stochastically) affect the quality of their submissions, and does not address the problem of optimally designing contests for agents who make only strategic participation choices, a difference that significantly alters the nature of agents’ incentives as discussed in §1. The model closest to ours from this literature is perhaps that of Taylor [26], where agents can make a random draw of output quality from a distribution F at a cost of c , similar to our model in §3. However, agents repeatedly make this choice over T periods (the duration of the contest), so that the central incentive problem faced by an agent in [26] is whether to make an *additional* draw at this cost. This repetition of the draw essentially mimicks a strategic choice of effort, albeit under different informational circumstances than in most other economic models of contests.

There is also a growing literature on modeling and analyzing incentives in online competitive environments such as crowdsourcing contests and online user-contributed content (see, e.g., [6, 1, 4, 8, 10, 3, 7, 15]). As in the contest design literature, almost all of this work models agents as making non-binary strategic effort (or quality) choices, and asks how to design incentives to elicit desired effort profiles from contestants. We note here that Cavallo and Jain [3] find—albeit in a completely different model from ours—that for specific distributions of agent abilities that satisfy a certain condition, the *efficient*, *i.e.*, first-best outcome (accounting for both the principal’s value and agents’ costs to effort), would be for some number of agents to exert maximum effort and for the remaining agents to exert no effort; Cavallo and Jain [3] then address the question of implementing the efficient outcome in equilibrium, *i.e.*, when agents strategically choose effort in response to the contest design, and prove both existence and impossibility results. While these ‘all or nothing’ effort choices bear an interesting resemblance to our model with agents who strategize only about participation, this resemblance is only superficial: the all-or-nothing choice in our model is an *assumption* about agents’ strategic behavior, while it arises as a structural property of efficient outcomes under certain conditions in [3].

The most closely related work to ours from this literature is [8] which investigates the implementability of optimal outcomes in widely used ‘best-contribution’ contests, in a model where agents with randomly drawn heterogeneous abilities make participation-only choices incurring a uniform cost. While that model corresponds precisely to the model in §3 of this paper, the nature of the questions addressed are fundamentally different—Ghosh and Hummel [8] *analyze* the implementability of optimal outcomes in a specific

contest structure *without* resource constraints; in contrast, we ask how to *design* the optimal contest under a given resource constraint, and identify which amongst several possible contest structures achieves the best outcome subject to this resource constraint.

We note that the model of simple agents who strategically reason only about participation choices is not new to this paper—in addition to [8], it appears at least in [9] in the context of privacy-sensitive users who make simple participation-only decisions about whether to share their data (although again in the context of equilibrium analysis rather than of optimal design), as well as possibly in several other settings in the vast economics literature. In the same spirit, we note that the general idea of moving from ideal to more realistic models is also not new to this paper: in a broad sense, our work belongs to a line of previous literature—albeit in completely different settings, with very different questions and models—that is driven by differences between the idealized models in traditional economic analyses and ‘real’ environments. This includes, for instance, the work on behavioral game theory [2], and building more realistic game-theoretic models of agent behavior as in [27], as well as designing mechanisms for more realistic assumptions about agents’ information or preferences, such as the literature on deliberative agents initiated in [18].

2 Model

We consider a model with n potential participants, each of whom strategically chooses only whether or not to participate in a given rank-order contest \mathcal{M} . An agent who decides to participate will produce an output of quality q_i , incurring a cost c_i , where agents’ quality-cost pairs are drawn from a known joint distribution F ; an agent makes her strategic participation choice based on whether her expected benefit from participation is large enough to cover her cost. Naturally, each agent’s choice depends on the participation choices of the remaining agents, since these determine the relative rank of her output q_i and therefore her prize. We will be interested in designing contests \mathcal{M} that maximize some increasing function of the qualities of submissions from the agents who choose to participate in the contest *in equilibrium*. We describe the model formally below.

Agents. There is a pool of n agents who are all potential participants in a contest. Agents are *simple*—they only strategically choose whether or not to *participate* in the contest, but do not strategically adjust their efforts, or equivalently the quality of their outputs, in response to the contest structure and environment. That is, *if* an agent decides to participate, she puts in a predetermined amount of effort incurring a corresponding cost, leading to some corresponding output quality; if her effort is too costly relative to her expected benefit from the resulting quality of output, she will simply not participate.

We model the fact that agents might have different abilities or skills relevant to the task posed by the contest via the quality of their output, as well as the cost they incur to produce this output—we denote the quality of the output that agent i can produce, if she chooses to participate, by q_i , and her associated cost by c_i . (An agent who does not participate produces no output, and incurs no cost.) An agent i is fully described by her quality-cost pair (q_i, c_i) ; we assume that the tuples (q_i, c_i) for each of the n agents are i.i.d. samples from a joint distribution $F(q, c)$ which is common knowledge to all agents and the contest designer. We additionally assume that the marginal distribution of qualities has no mass points: for every q_0 , the event $q = q_0$ has zero probability under F .

Contests. We suppose that the principal can observe the *rank-ordering* of the qualities q_i of participants’ submissions, and correspondingly consider rank-order mechanisms that award prizes to agents based (only) on the rank of their outputs. We use $\mathcal{M}(v_1, \dots, v_n)$ to denote the rank-order mechanism that awards a prize

v_j to the agent with the j^{th} rank, when submissions are arranged in order of decreasing quality with ties broken at random⁹. A mechanism $\mathcal{M}(v_1, \dots, v_n)$ is monotone if higher-ranked agents receive larger prizes, *i.e.*, if

$$v_1 \geq v_2 \geq \dots \geq v_n,$$

and monotone nonnegative if, in addition, $v_n \geq 0$. We will use the term contest to mean a monotone nonnegative rank-order mechanism henceforth in the paper.

Utility, strategies, and equilibrium. Agents are risk-neutral with quasi-linear utilities: an agent's utility, if she chooses to participate in a contest $\mathcal{M}(v_1, \dots, v_n)$, is the difference between her expected prize (determined by her rank, which depends on how her quality q_i ranks relative to the random draws $q_{i'}$ of the remaining agents who choose to participate) and her cost c_i to produce the output q_i . An agent who does not participate receives no prize and incurs no cost, and so obtains utility 0.

A strategy prescribes whether or not an agent with a given quality-cost draw should participate in a contest. Formally, a (pure) *strategy* π is a function mapping a quality-cost tuple (q, c) to a binary participation decision specifying whether or not to participate in a mechanism \mathcal{M} ; a player (q_i, c_i) following strategy π participates in \mathcal{M} if and only if $\pi(q_i, c_i) = 1$. A pure strategy profile (π_1, \dots, π_n) is a vector of strategies, one for each of the n players; a *symmetric* pure strategy profile is one where $\pi_i = \pi$ for all $i \in [n]$ and is denoted $\vec{\pi}$. Finally, a pure strategy profile in which $\pi_{i'} = \pi$ for all $i' \neq i$, but π_i may differ from π , will be denoted by $(\pi_i, \vec{\pi}_{-i})$.

Agents make their strategic participation choices based on the mechanism $\mathcal{M}(v_1, \dots, v_n)$ announced by the principal, the strategies $\vec{\pi}_{-i}$ of other agents, and their knowledge of the distribution F of quality-cost tuples (q, c) and the number of agents n . An agent i with quality and cost (q_i, c_i) decides whether to participate or not by comparing her expected utility from producing output q_i at cost c_i against the utility of 0 from non-participation; we assume that an agent who is indifferent between participation and non-participation always breaks ties in favor of participation¹⁰.

A symmetric pure-strategy profile $\vec{\pi}$ constitutes a Bayes-Nash equilibrium if no agent can profitably deviate from the recommendation of the strategy π given her draw (q_i, c_i) ; that is, (i) an agent with $\pi(q_i, c_i) = 1$, *i.e.*, who participates when playing according to π , obtains non-negative expected utility assuming that the remaining $n - 1$ agents with (q, c) drawn from F also play according to π , and (ii) the expected utility of an agent with $\pi(q_i, c_i) = 0$, if she were to participate while the remaining $n - 1$ agents play according to π , is negative. That is, no participant can beneficially deviate by not participating, and no non-participant would wish to participate.

For a given symmetric strategy profile $\vec{\pi}$, the set of participants in \mathcal{M} is a random variable constructed by choosing n i.i.d. draws of (q_i, c_i) from F , and then applying π to each of these (q_i, c_i) tuples. We use $S(\pi)$ to denote the random set of participants when all $n - 1$ agents follow the strategy profile π . Note that S depends only on the strategy profile chosen by agents, and does not depend at all on the mechanism \mathcal{M} .

Objective and constraints. We will be interested in the design of optimal contests, *i.e.*, an allocation of rewards (v_1, \dots, v_n) that maximizes some function of the qualities of submissions from the agents who

⁹Note that our assumption that the marginal distribution of q has no point masses implies that almost surely, the tie-breaking rule need not be invoked.

¹⁰In this paper, we always adopt assumptions on the joint distribution of (q_i, c_i) (such as no point masses) that are sufficient to ensure that agents only employ this tie-breaking rule with probability 0, so that this assumption does not affect our qualitative conclusions, but only makes our results easier to state.

choose to participate in a symmetric equilibrium¹¹ of the contest, subject to some constraint on the available rewards v_j . (Note that these rewards v_j may be either monetary, or non-monetary social-psychological rewards; see §1.) While there are various possible constraints on the rewards¹², we will focus in this paper on a sum constraint: $\sum_{j=1}^n v_j \leq V$. We assume throughout that the rewards, or prizes, v_j for each rank j are announced *prior* to the contest. Note this means that there might be leftover prize money for any particular instantiation of the agents' (q_i, c_i) pairs and strategic choices: depending on how many agents actually participate, the total sum of rewards given out to participating agents might be strictly less than V . That is, we consider rank-order mechanisms where the principal commits to a single allocation of V amongst the n ranks a priori, and cannot redistribute V (or equivalently, any leftover prize money) depending on the participation actually elicited, in contrast to hypothetical mechanisms where the principal might announce a full menu of prizes $((v_1^1), (v_1^2, v_2^2), \dots, (v_1^n, \dots, v_n^n))$ for each possible level of participation $k = 1, \dots, n$.

The outcome of a contest is determined by the output qualities q_i of agents who choose to participate in the contest. Let $\mathbf{q}(\pi) \in R^n$ denote the (random) vector consisting of the ordered output qualities of agents who participate when all agents play according to the strategy profile π . Formally, define an agent i 's output under π to be q_i if i participates (*i.e.*, $\pi(q_i, c_i) = 1$) and 0 if she does not ($\pi(q_i, c_i) = 0$); the (random) vector $\mathbf{q}(\pi)$ then holds the k^{th} largest of these n outputs $q_i \cdot \pi(q_i, c_i)$ in its k^{th} component.

The principal seeks to design mechanisms that maximize the expected value of some objective function f of the output qualities elicited in equilibrium. We define an optimal contest for a given objective f to be any mechanism $\mathcal{M}(v_1, \dots, v_n)$ such that \mathcal{M} has a symmetric pure-strategy equilibrium that maximizes $E[f(\mathbf{q}(\pi^*(\mathcal{M})))]$ ¹³ over all symmetric pure-strategy equilibria $\pi^*(\mathcal{M})$ of all feasible contests \mathcal{M} . Note that it is not clear, *a priori*, that there exists any mechanism with a symmetric pure-strategy equilibrium, nor that a given mechanism necessarily has such an equilibrium or that the equilibrium is unique when it exists. In Sections 3 and 4, we will prove existence and uniqueness of symmetric equilibria under the mild hypotheses that the distribution of q_i (in §3) and the joint distribution of (q_i, c_i) (in §4) are absolutely continuous with respect to Lebesgue measure, *i.e.*, they assign probability zero to every measure-zero set.

Throughout, we will be interested in objective functions that are *increasing* functions of the qualities of the elicited outputs: a function f is *increasing* if $f(\mathbf{q}_1) \geq f(\mathbf{q}_2)$ whenever the vector $\mathbf{q}_1 \succeq \mathbf{q}_2$. Informally, a function f of the set of elicited output qualities is increasing if higher-quality outputs are (weakly) preferred to lower-quality ones, and increasing participation to a superset of existing participants also makes the principal happier. We note that this assumption—that ‘more is better’, or equivalently, free disposal of output—is not entirely without loss of generality: it does not capture, for instance, search costs associated with sorting or finding better outputs in a larger set of outputs, or equivalently says that such costs are zero or negligible compared to the benefit of having a larger set of outputs. Nonetheless, the assumption of an increasing objective is reasonable in a vast number of settings, and essentially all the objective functions used in the (now vast) literature on contest design, including the two most commonly used objectives, namely the maximum output $\max_{i \in S} q_i$ and the sum of outputs $\sum_{i \in S} q_i$, are such increasing functions.

¹¹All contests that we consider in this paper have a symmetric equilibrium which is unique, and we restrict our attention to symmetric equilibria.

¹²For instance, if the v_i are attention rewards associated with content being displayed down a webpage, a natural constraint is a maximum per-slot attention constraint of the form $v_j \leq V_j$ [10].

¹³As before, the expectation is over random draws of (q_i, c_i) from F .

3 Homogenous participation costs

We begin by considering a special case of the distribution $F(q, c)$ where all agents have homogenous costs of participation, *i.e.*, $c_i = c$ for all i . A number of competitive environments can be reasonably described by a model with such homogenous participation costs—any contest where the participation cost c_i is primarily the cost of submission (either monetary, such as an entry fee, or the effort cost) rather than production (§1), or where the cost of producing an entry is largely independent of its quality (as in online content contribution settings where users differ in their abilities to contribute but incur similar contribution costs irrespective of quality, or in time-bound contests where the perceived cost is the time cost of participating in the contest), are well captured by a model with homogenous costs $c_i = c$.

We first analyze symmetric equilibria¹⁴ of contests when agents have homogenous costs in §3.1, and show that the equilibrium in any contest always consists of a threshold strategy. This means that the optimal contest for any increasing objective function of the designer is identical, and corresponds to a contest with the lowest threshold. We use this equilibrium analysis and an LP formulation of the contest design problem to prove our main result (§3.3), which is that optimal contests are *simple*: an optimal contest awards an equal prize to some number j^* of the highest-ranked contestants, and no prize to the rest. (Note that this structure is in contrast with the optimal contest structure in comparable models with strategic *effort* choices, where the optimal contest is either a winner-take-all contest or may award unequal prizes, depending on the curvature of agents’ effort cost functions [21].) The implications of this result for contest structures in various practical competitive environments are discussed in §1¹⁵. Finally, we investigate how the optimal number of prizes j^* , as well as the expected participation, vary with the ‘scale’ V/c of the contest in §3.4.

Throughout this section, we will use $F(q)$ to denote the marginal distribution from which agents’ qualities q_i are drawn. We will also assume that $\frac{V}{n} < c$, *i.e.*, there is not enough prize money to incentivize all agents to participate, since otherwise the optimal contest design problem is trivial: the contest with $v_j = \frac{V}{n} \geq c$ achieves full participation, since all agents receive a prize that covers or exceeds their cost to participation.

3.1 Equilibrium analysis

We first analyze equilibrium behavior in a general contest $\mathcal{M}(v_1, \dots, v_n)$ with simple agents who have homogenous costs. We show that with homogenous costs, the equilibria for any contest are threshold-strategy equilibria, where all agents with qualities q_i above some threshold participate, and the remaining agents do not. Before proving the result, we first introduce the function $c_{\mathcal{M}}(p)$ defined by

$$c_{\mathcal{M}}(p) = \sum_{j=1}^n v_j \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j}. \quad (1)$$

This function will turn out to be central to our analysis. Its value denotes the expected prize earned by a participant in the contest $\mathcal{M}(v_1, \dots, v_n)$ when each of the other $n-1$ agents has probability p of participating and placing ahead of that participant. Of course, the value of p depends on the type of that agent—her output quality q_i —as well as the strategies upon which each of the remaining $n-1$ contestants base their participation decisions, so that $p = p(q_i, \pi_{-i})$. When all agents play according to a common *threshold* strategy, though, if an agent with quality q_i participates then this probability of another agent participating

¹⁴Since we restrict attention to symmetric equilibria throughout this paper, we henceforth omit the modifier “symmetric” when discussing equilibria.

¹⁵See “Our contributions”.

and beating her is precisely $p = 1 - F(q_i)$, since i participates only if her only quality meets or exceeds the threshold.

The following lemma will be crucial to our analysis.

Lemma 3.1. *For any monotone rank-order mechanism $\mathcal{M} = \mathcal{M}(v_1, \dots, v_n)$, $c_{\mathcal{M}}(p)$ is non-increasing in p . If $v_1 \neq v_n$, then $c_{\mathcal{M}}(p)$ is strictly decreasing in p .*

Proof. As stated above, $c_{\mathcal{M}}(p)$ is equal to the expected prize gained by a contestant in contest \mathcal{M} with $n - 1$ other agents, when each of these agents (independently) places ahead of her with probability p . Intuitively, as we increase p , this contestant's (random) rank for larger values of p stochastically dominates the rank for smaller values of p , and the lemma follows since prize is a decreasing function of rank.

More rigorously, for $p_0 < p_1$, consider an urn containing three types of balls: red balls with probability p_0 , white balls with probability $p_1 - p_0$, and blue balls with probability $1 - p_1$. Consider a random experiment in which $n - 1$ balls are sampled independently (with replacement) from the urn, and j_0 is defined to be the number of red balls sampled while j_1 is defined to be the number of sampled balls that are either red or white. From the formula for $c_{\mathcal{M}}(p)$ one can see that $c_{\mathcal{M}}(p_0)$ is equal to the expected value of v_{j_0+1} while $c_{\mathcal{M}}(p_1)$ is equal to the expected value of v_{j_1+1} . Since $j_0 \leq j_1$ pointwise and v_1, \dots, v_n is a non-increasing sequence, it follows that $c_{\mathcal{M}}(p_0) \leq c_{\mathcal{M}}(p_1)$ as claimed. There is a positive probability that $n - 1$ white balls are drawn from the urn, and in the event that this happens $j_0 + 1 = 1$ while $j_1 + 1 = n$. Thus, assuming $v_1 > v_n$ we have $c_{\mathcal{M}}(p_0) < c_{\mathcal{M}}(p_1)$ which confirms the second claim in the lemma. \square

The next proposition asserts that rank-order mechanisms have unique equilibria in threshold strategies. The proof is deferred to Appendix A.

Proposition 3.1. *For any monotone rank-order mechanism \mathcal{M} , there is a symmetric equilibrium in threshold strategies where every agent with quality $q_i \geq q^*$ participates, and every agent with quality $q_i < q^*$ does not participate; the equilibrium threshold q^* satisfies $c_{\mathcal{M}}(1 - F(q^*)) = c$. This is the unique symmetric equilibrium of \mathcal{M} .*

3.2 Optimal contest design: LP formulation

The fact that there is a unique equilibrium in threshold strategies for every monotone rank-order mechanism $\mathcal{M} = \mathcal{M}(v_1, \dots, v_n)$ means that the (equilibrium) vector of outputs \mathbf{q}_1 in a mechanism \mathcal{M}_1 with threshold q_1^* dominates the vector of outputs \mathbf{q}_2 in a mechanism \mathcal{M}_2 with threshold $q_2^* > q_1^*$, for every possible realization of the qualities q_i : every agent who participates in equilibrium in the second mechanism also participates in the first, and there could possibly be additional agents who participate in \mathcal{M}_1 but not in \mathcal{M}_2 . This means that the problem of choosing the optimal mechanism for any objective $f(\mathbf{q})$ which is increasing in the vector of outputs \mathbf{q} becomes very simple: we want to find the mechanism with the allocation of rewards (v_1, \dots, v_n) that leads to the lowest equilibrium threshold $q^*(\mathcal{M})$.

Consider any mechanism $\mathcal{M}(v_1, \dots, v_n)$, and let $p = 1 - F(q^*)$ where q^* is the equilibrium threshold in $\mathcal{M}(v_1, \dots, v_n)$ (Proposition 3.1). Since all agents with qualities $q_i \geq q^*$ participate in \mathcal{M} in equilibrium, we can interpret p as the ex-ante probability (*i.e.*, prior to observing the draw of q_i) with which each agent will participate in the mechanism. We want to distribute the total prize V into (v_1, \dots, v_n) so as to maximize this participation probability $p = 1 - F(q^*)$, since maximizing this participation probability maximizes any increasing objective f of the principal running the contest.

Feasibility. First, note that the problem of *deciding* whether or not a *given* participation rate p is feasible for a given total prize V , and participation cost c , can be written as a feasibility LP in the variables (v_1, \dots, v_n) :

$$\begin{aligned} v_1 &\geq v_2 \geq \dots \geq v_n \geq 0, \\ \sum_{j=1}^n v_j &\leq V, \\ \sum_{j=1}^n v_j \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j} &= c. \end{aligned} \tag{2}$$

The first two constraints say that we restrict ourselves to monotone nonnegative mechanisms \mathcal{M} that distribute a total prize of at most V . The third constraint is the *incentive* constraint; it accounts for the fact that agents make strategic (albeit simple) choices in response to the incentives created by the mechanism $\mathcal{M}(v_1, \dots, v_n)$, and ensures that $p = 1 - F(q^*)$ actually describes the equilibrium threshold in \mathcal{M} .

If this LP is feasible, there exists an allocation of the total reward V into rewards (v_1, \dots, v_n) that elicit an equilibrium participation probability of p ; if not, no contest with total available prize V can achieve a participation rate of p when agents have cost c to participation.

Optimality. The feasibility LP asks *whether* it is possible to achieve some level of participation p given a total reward V , when agents have participation costs c . The question we would actually like to ask is *what* is the highest participation rate we can achieve given the pair (V, c) , and how—*i.e.*, using what allocation scheme (v_1, \dots, v_n) ? Since the incentive constraint—the third constraint in (2)—is nonlinear in p , we cannot write the problem of maximizing p over all allowable choices of (v_1, \dots, v_n) for given values of c and V as an LP. However, the monotonicity between c and p established in Lemma 3.1 will allow us to approach this question via an LP anyway.

Fix a mechanism $\mathcal{M} = \mathcal{M}(v_1, \dots, v_n)$: for every p , the function $c_{\mathcal{M}}(p)$ defined in (1) gives the precise participation cost at which $\mathcal{M}(v_1, \dots, v_n)$ has equilibrium participation rate p . Now for a *given* p , consider the following LP in the variables (v_1, \dots, v_n) :

$$\begin{aligned} \text{maximize} \quad & \sum_{j=1}^n v_j \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j} \\ \text{s.t.} \quad & v_1 \geq v_2 \geq \dots \geq v_n \geq 0, \\ & \sum_{j=1}^n v_j = V. \end{aligned} \tag{3}$$

Note that the objective of this LP is precisely $c_{\mathcal{M}}(p)$. Therefore, for a given $p \in [0, 1]$, the solution to this optimization problem gives the largest participation cost c at which a participation rate of p can be supported in equilibrium by optimizing over how the prize V is distributed across ranks.

Lemma 3.1, which asserts the monotonicity of $c_{\mathcal{M}}(p)$, allows us to use this LP to address the question of *optimal* contests, *i.e.*, contests that achieve the best possible participation rate p for a given total prize V and participation cost c , as follows. Define the function $c^*(p)$ to be the optimal value of the LP (3) for each $p \in [0, 1]$: $c^*(p) = \max_{\mathcal{M}(v_1, \dots, v_n)} c_{\mathcal{M}}(p)$ describes the maximum participation cost c at which a participation rate of p can still be supported in equilibrium by appropriately choosing the most favorable reward allocation (v_1, \dots, v_n) . From Lemma 3.1, the function $c^*(p)$ is strictly decreasing, since it is the pointwise maximum of strictly decreasing functions. (Our assumption that $\frac{V}{n} < c$ implies that $v_1 \neq v_n$ for any contest that elicits non-zero participation, since if $v_1 = v_n$, the prize $v_j \leq \frac{V}{n}$ for all ranks j so that no agent participates). Hence $c^*(p)$ is also invertible; define the inverse of this function to be $p^*(c)$. The function $p^*(c)$ is the highest participation rate that can be supported in equilibrium with a total prize V when agents have participation cost c , and this rate is achieved by any contest that maximizes the objective in (3) for $p = p^*(c)$.

Note that the invertibility of $c^*(p)$ is important *not* because we cannot answer the question of the best achievable p^* for a given (V, c) without it, but rather because it allows us to say that an allocation

(v_1, \dots, v_n) achieving the optimum of (3) for any specified p is indeed an optimal contest (for the value of c corresponding to the LP optimum). Without invertibility, we could still define $p^*(c)$ as the largest p for which $c^*(p)$ was equal to c , to answer the question of the best achievable p for a given (V, c) pair. However, studying the optimal solutions to the LP (3) would not necessarily give us any insights into the *structure* of optimal contests, since we are interested in understanding what contest designs $\mathcal{M}(v_1, \dots, v_n)$ achieve the best possible outcome p for given cost c , rather than the contests that have the lowest c for a given p .

3.3 Optimal contests are simple

We will now use the LP formulation (3) to derive the structure of optimal contests. Our main result is that optimal contests are *simple*: for every value of total prize V and participation cost c , there is an optimal contest which awards j equal prizes of value V/j each to some top j ranks, and a prize of 0 to all ranks $j + 1$ and below. To prove this, we reformulate the LP (3) by changing variables. For $j = 1, \dots, n$, let $w_j = j(v_j - v_{j+1})$, where we define $v_{n+1} = 0$. The constraint on the total allocated prize can be written as

$$V = \sum_{j=1}^n v_j = (v_1 - v_2) + 2(v_2 - v_3) + 3(v_3 - v_4) + \dots + n(v_n - v_{n+1}) = w_1 + \dots + w_n. \quad (4)$$

The nonnegativity and monotonicity constraints on the prizes v_j , that is, $v_j \geq v_{j+1}$ for $j = 1, \dots, n-1$ and $v_n \geq 0$, simply translate to $w_j \geq 0$ for $j = 1, \dots, n$. Finally, we rewrite the objective in terms of the w_j :

$$\begin{aligned} \sum_{j=1}^n v_j \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j} &= \sum_{j=1}^n \left(\sum_{k=j}^n \frac{w_k}{k} \right) \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j} \\ &= \sum_{j=1}^n w_j \left(\frac{1}{j} \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \right). \end{aligned}$$

We can therefore rewrite the LP (3) in the variables (w_1, \dots, w_n) as follows:

$$\begin{aligned} \text{maximize}_{\{w_j\}} \quad & \sum_{j=1}^n w_j \left(\frac{1}{j} \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \right) \\ \text{s.t.} \quad & \sum_{j=1}^n w_j = V, \\ & w_j \geq 0, \quad j = 1, \dots, n. \end{aligned} \quad (5)$$

An optimal solution to this LP is set $w_{j^*} = V$ for the index with the largest coefficient multiplying w_j in the objective, *i.e.*, the w_j corresponding to

$$j^* = \arg \max_{1 \leq j \leq n} \left\{ \frac{1}{j} \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \right\}, \quad (6)$$

and $w_j = 0$ for all $j \neq j^*$. We note here that there are only finitely many values of p for which the arg max is non-unique, because if the maximum is achieved at two different values j, j' then this implies that p satisfies a non-trivial polynomial equation, *i.e.*

$$\frac{1}{j} \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} = \frac{1}{j'} \sum_{\ell=1}^{j'} \binom{n-1}{\ell-1} p^{\ell-1} (1-p)^{n-\ell}.$$

What contest structure, *i.e.*, actual awarded prizes (v_1, \dots, v_n) does this LP solution correspond to? Recall our transformation of variables $w_j = j(v_j - v_{j+1})$. With $w_{j^*} = V$ and the remaining w_j set to 0, the prizes v_j in an optimal contest are then:

$$v_j = \frac{V}{j^*}, \quad 1 \leq j \leq j^*, \quad \text{and} \quad v_j = 0, \quad j^* < j \leq n.$$

Recall that the equation (6) defining j^* has a unique solution apart from a finite number of exceptional values of p . These exceptions correspond, under the strictly monotone function $c^*(p)$, to a finite number of exceptional values of c ; for all other c there is a unique optimal contest and it is a simple contest. We summarize this result in the following theorem.

Theorem 3.1. *For every value of total prize V and participation cost c , there is an optimal contest which awards equal prizes V/j^* to some top j^* ranks, and a prize of 0 to all ranks $j^* + 1$ and below. Apart from finitely many exceptional values of c , this is the unique optimal contest.*

3.4 Understanding the behavior of the optimal contest

The result in Theorem 3.1 illustrates the *structure* of optimal contests, namely that there are only two levels of prizes—equal prizes to some top j^* contestants, and a prize of 0 to all lower-ranked contestants. This description of the optimal contest is complete except for one crucial piece of information: what is the value of j^* ? In this subsection, we seek to understand how the total number of prizes in the optimal contest, j^* , varies with the parameters of the contest: the total prize V , participation cost c , and number of potential contestants, n .

A basic question regarding the optimal number of prizes, j^* , concerns comparative statics: is j^* monotonic as a function of participation cost? It is intuitively clear that as the cost of participation, c , increases, the *ex ante* probability of an individual agent participating in the equilibrium of the optimal contest, $p^*(c)$, must decrease; this intuition was confirmed rigorously in §3.2. On the other hand, the monotonicity of j^* as a function of c is much less apparent: as the cost of participation increases, should the contest designer offset the cost increase by increasing the magnitude of prizes (which necessitates reducing the total number of prizes offered, and thus the probability of winning a prize) or by increasing the total number of prizes (which increases the probability of winning a prize but reduces the magnitude of each prize)? The following proposition implies, among other things, that an optimal contest designer responds to increasing participation costs by offering the same or smaller number of prizes. All proofs in this section are deferred to Appendix ??.

Proposition 3.2. *Fix a total prize, V , and number of potential contestants, n . Let \mathcal{M}^j denote a simple contest that awards equal prizes V/j to the top j ranks.*

1. *There is a decreasing sequence $V = c_1 > c_2 > \dots > c_n > c_{n+1} = 0$, such that for any participation cost $c < V$, the simple contest \mathcal{M}^j is optimal if and only if $c_{j+1} \leq c \leq c_j$.*
2. *For $1 < j \leq n$, the equation $c_{\mathcal{M}^j}(p) = c_{\mathcal{M}^{j-1}}(p)$ has a unique solution p_j in the interval $(0, 1)$, and $c_j = c_{\mathcal{M}^j}(p_j) = c_{\mathcal{M}^{j-1}}(p_j)$.*

Proposition 3.2 allows us to draw a number of quantitative conclusions about j^* , the number of prizes offered by the optimal contest. The following proposition identifies a precise criterion for optimality of the winner-take-all contest \mathcal{M}^1 ; for every fixed j it identifies an asymptotic formula for p_j as $n \rightarrow \infty$, yielding a criterion for optimality of \mathcal{M}^j that holds for all sufficiently large population sizes; and finally, it presents an asymptotic estimate of the number of prizes offered by the optimal contest, and the expected number of

participants in equilibrium, in terms of the ratio $\nu = V/c$. We call this ratio the *scale* of the contest, because ν is clearly an upper bound on the number of prizes that can be awarded, if each prize must be at least large enough to cover agents' participation costs. Proposition 3.3 shows that both the number of prizes awarded and the expected number of participants are asymptotic to ν minus a lower-order term. In other words, as the total prize offered grows large compared each individual agent's participation cost, the equilibrium of the optimal contest asymptotically approaches the ideal of "perfect coordination", in which each of the top $\lfloor V/c \rfloor$ agents (ranked by quality) is asked to participate and paid just enough to cover their participation cost, despite the total lack of coordination among agents.

Proposition 3.3. *Fix V, c , and let $j^{*,n}$ denote the number of prizes offered by the optimal contest with population size n . Let $p^{*,n}$ denote the probability that an individual agent participates in the equilibrium of that contest, and let $\lambda^{*,n} = p^{*,n} \cdot n$ denote the expected number of participants in that equilibrium.*

1. *The winner-take-all contest is optimal if and only if $\frac{V}{c} \leq \left(1 + \frac{1}{n-1}\right)^{n-1}$. In particular, for $\frac{V}{c} > e$ a winner-take-all contest is never optimal, whereas for $\frac{V}{c} < e$ it is optimal whenever the population size, n , is sufficiently large.*
2. *As $n \rightarrow \infty$, the sequence $\lambda^{*,n}$ converges to the unique $\lambda^* > 0$ satisfying*

$$\max_{1 \leq j \leq V/c} \left\{ \frac{V}{j} \sum_{k=0}^{j-1} \frac{e^{-\lambda} \lambda^k}{k!} \right\} = c. \quad (7)$$

Excluding a countable set of values for c , the maximum in (7) is achieved at a unique value j^ , and the sequence $j^{*,n}$ converges to j^* as $n \rightarrow \infty$.*

3. *Letting $\nu = V/c$, for any $n > 2\nu$ the values of $\lambda^{*,n}$ and $j^{*,n}$ satisfy*

$$\lambda^{*,n} = \nu - \Theta(\sqrt{\nu \log \nu}), \quad j^{*,n} = \nu - \Theta(\sqrt{\nu / \log \nu}). \quad (8)$$

4 Arbitrary quality-cost distributions

We now proceed to settings where agents have heterogeneous costs c_i in addition to differing output qualities q_i , modeled via a general joint distribution $F(q, c)$ on agents' costs and output qualities. Throughout this section, we assume that $F(q, c)$ is absolutely continuous with respect to Lebesgue measure, *i.e.*, every measure-zero set of (q, c) pairs also has zero probability under the distribution $F(q, c)$.

When output qualities and costs can have arbitrary correlations, much of the structure we saw in the previous section vanishes. Specifically, equilibria no longer need possess a threshold structure, where all agents whose output quality is greater than a threshold participate and the remaining agents do not. (For example, a very high-ability agent producing very high quality might also have a prohibitively high cost preventing entry, whereas a lower-ability agent might still find it profitable to participate because of an adequately low participation cost c_i .) The absence of such a threshold equilibrium structure under heterogeneous costs means that unlike with homogeneous costs, there may be no single contest \mathcal{M} that is optimal for all increasing objective functions f . (See §4.3 for a simple example that illustrates this point; we present a distribution for which there is no contest \mathcal{M} that simultaneously optimizes the sum objective and the max objective.) Specifying an optimal contest, therefore, involves also choosing the objective that the mechanism designer wants to maximize.

Motivated by the fact that many contests are aimed at obtaining one best innovation or solution to a problem, or one best design, and so on, a significant fraction of the literature on contest design has focused on the max objective, *i.e.*, the expected quality of the highest-ranked submission. In keeping with this motivation and literature, we focus on the max objective in this section (although we note that a number of other objectives, such as the sum or more generally, the sum of the top k outputs, have also been used in prior work, and are potentially well-motivated in the settings we consider as well). Given a total prize V , and n agents with types (q_i, c_i) drawn from F , we will therefore define an *optimal contest* to mean any mechanism $\mathcal{M} = \mathcal{M}(v_1, \dots, v_n)$ such that some symmetric pure-strategy equilibrium of \mathcal{M} achieves the maximum $E[\max\{q_i \mid i \in \mathcal{S}(\pi^*(\mathcal{M}))\}]$ over all symmetric pure-strategy equilibria $\pi^*(\mathcal{M})$ of all contests \mathcal{M} with $\sum v_j = V$, where the set $\mathcal{S}(\pi)$ is as defined in §2.¹⁶

With general type distributions F , an agent’s type is necessarily two-dimensional and cannot be characterized by a single number. This is in contrast with §3, where agents had homogeneous costs (an agent’s type is essentially just her output quality q_i), as well as models with strategic effort choices in the prior literature where an agent’s type is typically either her ability, or her cost per unit effort. With such arbitrary two-dimensional types, an equilibrium analysis that characterizes equilibrium outcomes in an arbitrary rank-order mechanism appears to be fairly challenging, which means that the standard approach to solving optimal contest design problems—first explicitly solve for equilibrium outcomes for general rank-order mechanisms $\mathcal{M}(v_1, \dots, v_n)$, and then optimize over the choice of feasible (v_1, \dots, v_n) to maximize the objective achieved by a mechanism in equilibrium—cannot be directly applied here.

Rather than derive the optimal mechanism, therefore, we ask how well *simple* mechanisms — mechanisms that award equal prizes V/j to the highest-ranked j agents, for some j — can perform, in terms of the expectation of the highest-quality output in equilibrium relative to an optimal mechanism. In this section, we will show that the winner-take-all mechanism (corresponding to $j = 1$), which awards the entire prize budget V to the participant with the highest quality output, is a 3-approximation to the optimal mechanism.

Existence and uniqueness of pure strategy equilibria Before delving into the analysis of symmetric equilibria of contests under arbitrary joint distributions F , we pause here to mention some basic facts which ensure that the analysis is meaningful. Proofs of these facts are given in Appendix B. First, for every contest \mathcal{M} there exists a symmetric mixed-strategy equilibrium. Second, all mixed-strategy equilibria of \mathcal{M} are “pure almost everywhere,” in the sense that the set of types that randomize in equilibrium has measure zero. Third, the symmetric equilibrium of \mathcal{M} is unique up to an indeterminacy on a measure-zero set of types (q, c) .

4.1 Sub-equilibrium lemma

As discussed above, solving for the equilibrium of a mechanism under arbitrary joint distributions F is challenging. To circumvent the need to solve for equilibria, we instead relate the expected outcome in an equilibrium of the winner-take-all mechanism \mathcal{M}_{WTA} to that in an optimal mechanism by identifying and analyzing *sub-equilibrium* strategy profiles in \mathcal{M}_{WTA} , as defined below.

Definition 4.1. A strategy profile (π_1, \dots, π_n) is a sub-equilibrium of a given contest if it satisfies *ex interim individual rationality*: for any player i and type (q_i, c_i) such that $\pi_i(q_i, c_i) = 1$, the expected payoff of player i from participating, conditional on having type (q_i, c_i) , is greater than or equal to c_i when $n - 1$ remaining players with types drawn from $F(q, c)$ play according to π . (Note that individual rationality

¹⁶As before, we restrict our attention to symmetric equilibria and henceforth omit the modifier “symmetric.”

vacuously holds for non-participants, i.e., agents with $\pi_i(q_i, c_i) = 0$, since non-participation yields zero utility and is therefore always individually rational.) We say that a pure strategy π is a sub-equilibrium if the corresponding symmetric strategy profile $\bar{\pi}$ with $\pi_i = \pi$ for all i is a sub-equilibrium.

A sub-equilibrium is a strategy profile that is equilibrium-like, but not quite an equilibrium: all agents who do participate derive non-negative utility, so that their action choice is indeed a best response, but there might be non-participating agents who would derive non-negative expected utility from participation given other agents' strategies, so that these agents' choices do not constitute a best response. A sub-equilibrium is therefore a strategy profile with 'too little' participation, where all *participants* are playing a best response, but non-participants might have a profitable deviation.

The interpretation of sub-equilibria as strategy profiles with 'too little' participation suggests that if π^S is a sub-equilibrium and π^E is an equilibrium, then the quality of an agent's output (treating non-participation as zero quality) in π^E should stochastically dominate the agent's output quality in π^S . Lemma 4.1 below confirms this intuition. Note, however, that it is not obvious that this intuition should be valid: some agent types who choose non-participation in π^S may choose to participate in π^E , but the participation of these types may adequately decrease the probability of winning for other agent types to cause them to flip their choice from participation to non-participation, as a rational response to the increased competition. That these effects should lead, on net, to stochastic dominance appears *a priori* to be far from evident.

Lemma 4.1 (Sub-equilibrium Lemma). *Fix a joint distribution $F(q, c)$ of types in an n -player contest $\mathcal{M} = \mathcal{M}(v_1, \dots, v_n)$. Assume that $v_1 \neq v_n$ and that the marginal distribution of q has no point masses. If pure strategy π^S is any sub-equilibrium and pure strategy π^E is any equilibrium of \mathcal{M} , then the random variable $q \cdot \pi^E(q, c)$ stochastically dominates the random variable $q \cdot \pi^S(q, c)$.*

Proof. Let F_S, F_E denote the cumulative distribution functions of the random variables $q \cdot \pi^S(q, c)$ and $q \cdot \pi^E(q, c)$, respectively. The lemma asserts that $F_S(x) \geq F_E(x)$ for all x . Our assumption that the marginal distribution of q has no point masses implies that the marginal distributions of $q \cdot \pi^S(q, c)$ and $q \cdot \pi^E(q, c)$ likewise have no point masses and hence F_S, F_E are continuous functions on the set of non-negative real numbers, \mathbb{R}_+ .

If an agent i with type (q, c) chooses to enter the contest \mathcal{M} , and all other agents play strategy profile π^S_{-i} , then the probability that a given one of those opponents participates and ranks above agent i is $1 - F_S(q)$. Therefore, if agent i chooses to participate, the probability that she ranks j^{th} is equal to $\binom{n-1}{j-1} (1 - F_S(q))^{j-1} F_S(q)^{n-j}$ and her expected prize is equal to $\sum_{j=1}^n v_j \binom{n-1}{j-1} (1 - F_S(q))^{j-1} F_S(q)^{n-j}$. Recalling the definition of the function $c_{\mathcal{M}}$ from §3, we see that an agent with type (q, c) who chooses to participate when the opponents play strategy profile π^S_{-i} will gain an expected reward of $c_{\mathcal{M}}(1 - F_S(q))$.

To prove that $F_S(x) \geq F_E(x)$ for all x , we argue by contradiction. Assume that $F_S(x_0) < F_E(x_0)$. Since the function $G = F_E - F_S$ is continuous on \mathbb{R}_+ , there is an open interval I containing x_0 such that G is strictly positive on $I \cap \mathbb{R}_+$. The union of all such open intervals is itself a non-empty open interval (\underline{x}, \bar{x}) ; here we allow $\underline{x} = -\infty$ or $\bar{x} = +\infty$, or both. We know that $\bar{x} > 0$ because $\bar{x} > x_0$. In $0 < \bar{x} < \infty$, then the continuity of G implies that $G(\bar{x}) = 0$ and that

$$\lim_{x \rightarrow \bar{x}} G(x) = 0. \quad (9)$$

If $\bar{x} = \infty$ then $F_E(x)$ and $F_S(x)$ both converge to 1 as $x \rightarrow \bar{x}$, and so again (9) holds.

For any $(q, c) \in \mathbb{R}_+^2$ such that $\pi^S(q, c) = 1$ and $\underline{x} < q < \bar{x}$, we have

$$c \leq c_{\mathcal{M}}(1 - F_S(q)) < c_{\mathcal{M}}(1 - F_E(q)), \quad (10)$$

where the first inequality expresses the individual rationality constraint in the definition of sub-equilibrium, and the second inequality holds because $F_E - F_S$ is positive on interval (\underline{x}, \bar{x}) and the function $c_{\mathcal{M}}(1 - y)$ is strictly increasing in y (by Lemma 3.1). Since π^E is an equilibrium, inequality (10) implies that $\pi^E(q, c) = 1$, since $c_{\mathcal{M}}(1 - F_E(q))$ is the expected prize to an agent with quality q in the strategy profile π^E . Thus, we have shown that $\pi^S(q, c) = 1$ implies $\pi^E(q, c) = 1$ when $\underline{x} < q < \bar{x}$. Now, if $\underline{x} < x_0 < x_1 < \bar{x}$, we have

$$\begin{aligned} F_S(x_1) - F_S(x_0) &= \Pr_{(q,c)}(x_0 < q < x_1, \pi^S(q, c) = 1) \\ &\leq \Pr_{(q,c)}(x_0 < q < x_1, \pi^E(q, c) = 1) = F_E(x_1) - F_E(x_0) \end{aligned}$$

and by rearranging terms on the left and right sides we find that $G(x) \geq G(x_0) > 0$ for all $x \in (x_0, \bar{x})$. This means that $\lim_{x \rightarrow \bar{x}} G(x)$ cannot be 0, contradicting (9) and completing the proof. \square

Recall from §2 that $\mathcal{S}(\pi)$ is the (random) set of participants when all agents make participation decisions according to strategy profile π . The following immediate corollary of Lemma 4.1 will be needed in the sequel.

Corollary 4.1. *Fix a joint distribution $F(q, c)$ of types in an n -player winner-take-all contest \mathcal{M}_{WTA} . Assume that the marginal distribution of q has no point masses. If pure strategy π^S is any sub-equilibrium and pure strategy π^E is any equilibrium of \mathcal{M} , then $\max\{q_i | i \in \mathcal{S}(\pi^E)\}$ stochastically dominates the random variable $\max\{q_i | i \in \mathcal{S}(\pi^S)\}$.*

4.2 Winner-take-all is a 3-approximation of the optimal contest

We will now sketch the proof that the expected value of the maximum quality entry in any pure-strategy equilibrium of the winner-take-all contest \mathcal{M}_{WTA} is no smaller than one-third of the expected maximum quality in the best equilibrium of an optimal contest \mathcal{M}^* .

Consider an optimal contest \mathcal{M}^* and let π^E denote its symmetric equilibrium (recall that π^E is unique). We use $\mathcal{E}^* = \mathcal{S}(\pi^E(\mathcal{M}^*))$ to denote the random set of participant types in the equilibrium strategy profile $\pi^E(\mathcal{M}^*)$ of \mathcal{M}^* , and let the random variable $m(\mathcal{E}^*) = \max\{q_i | i \in \mathcal{E}^*\}$ denote the maximum quality in a particular instantiation of the random set \mathcal{E}^* (recall that the principal's objective is to maximize the expected highest quality). Note that

$$m(\mathcal{E}^*) \leq \max\{q_i | i \in \mathcal{E}^*; c_i \leq \frac{V}{2}\} + \max\{q_i | i \in \mathcal{E}^*; c_i > \frac{V}{2}\},$$

since the highest-ability agent that participates in \mathcal{M}^* for any given random instantiation of \mathcal{E}^* has cost either less than, equal to, or greater than $V/2$. (We define the max of an empty set to be zero.)

We will use OPT to denote the expected value of $m(\mathcal{E}^*)$ over the random draws of the n types (q, c) . Since the inequality above holds pointwise, it also holds in expectation:

$$\text{OPT} = E[m(\mathcal{E}^*)] \leq E\left[\max\{q_i | i \in \mathcal{E}^*; c_i \leq \frac{V}{2}\}\right] + E\left[\max\{q_i | i \in \mathcal{E}^*; c_i > \frac{V}{2}\}\right]. \quad (11)$$

Now consider the winner-take-all mechanism \mathcal{M}_{WTA} . Denote the expected value of the maximum output q_i amongst the set of agents participating in the equilibrium $\pi^E(\mathcal{M}_{\text{WTA}})$ of the winner-take-all mechanism by WTA, i.e.,

$$\text{WTA} = E[\max\{q_i | i \in \mathcal{S}(\pi^E(\mathcal{M}_{\text{WTA}}))\}].$$

We will show that $3 \cdot \text{WTA} \geq \text{OPT}$ by showing that the first term in (11) is no larger than twice WTA and the second term in (11) is no larger than WTA. That is, imagine the expected quality of the best submission

in the optimal contest as ‘split’ across two subpopulations—the low-cost subpopulation (with $c_i \leq V/2$), and the high-cost subpopulation (with $c_i > V/2$). We show that the winner-take-all mechanism does at least as well as the optimal mechanism ‘restricted’ to the high-cost subpopulation, and no worse than half as badly as the optimal mechanism ‘restricted’ to the low-cost subpopulation. Note that the nature of these comparisons should not be entirely unexpected from our results in §3—in a homogeneous-cost model, the winner-take-all mechanism is suboptimal when participation costs are small relative to the total prize, but optimal when the cost c is large enough compared to V , suggesting that \mathcal{M}_{WTA} is likely to ‘lose’ relative to the performance of an optimal contest with a low-cost subpopulation, but not a high-cost subpopulation. The following two inequalities summarize this discussion.

$$\text{WTA} \geq \frac{1}{2} E[\max\{q_i \mid i \in \mathcal{E}^*; c_i \leq \frac{V}{2}\}] \quad (12)$$

$$\text{WTA} \geq E[\max\{q_i \mid i \in \mathcal{E}^*; c_i > \frac{V}{2}\}] \quad (13)$$

The proofs of both inequalities are presented in Appendix C. Both proofs rely on Lemma 4.1, the sub-equilibrium lemma: in both cases, the proof proceeds by constructing a sub-equilibrium of \mathcal{M}_{WTA} and showing that the expectation of the maximum quality in that sub-equilibrium—for brevity, the *value* of the sub-equilibrium—is an upper bound on the corresponding term in (11), or half of that term in the case of Inequality (12). To prove Inequality (12), we use a sub-equilibrium in which agents participate if their cost is at most $\frac{V}{2}$ and their quality is at least μ , where μ is chosen so that the probability that $\max\{q_i \mid c_i \leq \frac{V}{2}\}$ exceeds μ is exactly $\frac{1}{2}$. This is easily seen to be a sub-equilibrium: agents only participate if their probability of winning is at least $\frac{1}{2}$ in which case their expected prize is at least $\frac{V}{2}$, while their cost is at most $\frac{V}{2}$, which guarantees a non-negative utility. The proof that the value of the sub-equilibrium is at least the RHS of Inequality (12) is similar to Samuel-Cahn’s [1984] proof of the Prophet Inequality. To prove Inequality (13) we use a different sub-equilibrium, where the set of participating agent types equals the set of all types that participate in the equilibrium of the optimal contest *and* have cost greater than $\frac{V}{2}$. The value of this strategy profile is precisely the second term in (11); the difficulty here is to prove that the strategy profile is indeed a sub-equilibrium. We do this by showing that any monotone rank-order mechanism is payoff equivalent to a randomization over simple contests, *i.e.*, contests which award j equal prizes. This means that for a type (q, c) that participates in the equilibrium of the optimal contest, it must be individually rational to participate in at least one simple contest. An agent with cost greater than $\frac{V}{2}$ cannot find it individually rational to participate in any simple contest that awards more than one prize, so for the participating types in our putative sub-equilibrium, participation in a simple contest with one prize—that is, a winner-take-all contest—must be individually rational. This shows that the putative sub-equilibrium is actually a sub-equilibrium and completes the proof.

We note that neither of these proofs rely on any structural property specific to the optimal contest: all steps in both proofs remain valid if we instead compare \mathcal{M}_{WTA} to an arbitrary monotone rank-order mechanism respecting the sum constraint ($\sum_{j=1}^n v_j = V$), whether or not it is optimal. In fact, this aspect of our proofs is in some sense inevitable since we have almost no insight into the structure of the optimal contest in this setting with two-dimensional agent types resulting from arbitrary quality-cost distributions; the analysis therefore necessarily relies on properties that hold for all feasible contests.

Inequalities (12) and (13) combine to yield our main approximation result.

Theorem 4.1. *The expected maximum output in the winner-take-all mechanism is at least one-third that in the optimal mechanism: $\text{OPT} \leq 3 \cdot \text{WTA}$.*

4.3 Dependence on the objective function

We conclude this section with a simple example illustrating an important difference between the general two-parameter setting studied in this section and the homogeneous-cost setting in §3. Unlike the case of homogeneous costs, for a general joint distribution of output quality and cost, there may be no single contest that is optimal for all choices of the increasing objective function f : Example 4.1 below presents a joint distribution $F(q, c)$ such that the optimal contest for the sum objective differs from the optimal contest for the max objective.

Example 4.1. *Suppose that the joint distribution $F(q, c)$ is supported on the squares $[1, 1 + \epsilon] \times [1 - \epsilon, 1]$ and $[20, 21] \times [\frac{9V}{10} - 1, \frac{9V}{10}]$, with probability density $1/(2\epsilon^2)$ on the first square and $1/2$ on the second one. (Thus, to sample from $F(q, c)$ one first selects among the two squares with equal probability, then one chooses a uniformly random point in the selected square.) Denote the types (q, c) with $q > 1 + \epsilon$ as high types and those with $q \leq 1 + \epsilon$ as low types. Assume that $0 < \epsilon \ll 1$ and $n \gg V \gg 1$.*

Choose $x \in [20, 21]$ to satisfy $\Pr_{(q,c) \sim F}(q < x) = (9/10)^{1/n}$. Equivalently,

$$\Pr(\max\{q_1, \dots, q_n\} < x) = \frac{9}{10}$$

when n independent types (q_i, c_i) are sampled from F . An agent whose output quality is greater than x has probability at least $9/10$ of having the highest ability among all agents. Such an agent gains an expected prize of at least $9V/10$ in a winner-take-all contest \mathcal{M}_{WTA} , no matter what strategy profile the other agents use, and therefore the equilibrium of \mathcal{M}_{WTA} has full participation of agents with quality greater than x . The probability that at least one of the n agents samples a type (q_i, c_i) with $q_i > x$ is $\frac{1}{10}$, so the expectation of the maximum output quality in the equilibrium of the winner-take-all contest is at least $\frac{x}{10} > 2$.

In comparison, any contest with $v_1 < \frac{9V}{10} - 1$ cannot elicit the participation of any high-type agents, and hence the expectation of the max objective, in the equilibrium of such a contest, is less than $1 + \epsilon$. This shows that the optimal contest for the max objective must satisfy $v_1 \geq \frac{9V}{10} - 1$.

Now consider the sum objective. If \mathcal{M} denotes a contest that offers $V/2$ equal prizes of size 2, we claim that in the equilibrium of \mathcal{M} , the expected sum of output qualities is at least $V/4$. Indeed, high types do not participate in equilibrium (the prize is less than their cost) so only the low types participate. Let $y \in [1, 1 + \epsilon]$ be such that, in expectation, exactly $V/4$ of the low types among $\{(q_1, c_1), \dots, (q_n, c_n)\}$ satisfy $q_i > y$. By Markov's Inequality, the probability that more than $V/2$ low types with $q_i > y$ participate is at most $1/2$, and hence an agent i with $q_i > y$ has probability at least $1/2$ of winning a prize of size 2. If the agent is of low type, then the expected prize exceeds the agent's cost, and consequently all of the low-type agents with $q_i > y$ will participate in equilibrium. The expected number of such agents is $V/4$ and each of them contributes an output of quality at least 1, so that expectation of the sum objective in equilibrium is at least $V/4$ as claimed.

In comparison, in any contest with $v_1 \geq \frac{9V}{10} - 1$, the total prize awarded to contestants in ranks 2 and higher is at most $\frac{V}{10} + 1$. Define a random variable Z to be equal to the total cost expended by the agents who participate and place in ranks 2 and higher. Individual rationality dictates that $E[Z] \leq \frac{V}{10} + 1$. However, the ratio q/c is at most $\frac{1+\epsilon}{1-\epsilon}$ for all points (q, c) in the support of the type distribution. Consequently, the expected total quality of the submissions from agents who place in ranks 2 and higher is bounded above by $\frac{1+\epsilon}{1-\epsilon} (\frac{V}{10} + 1)$. The top-ranked agent has output quality at most 21, and hence the expectation of the sum objective in equilibrium is at most $21 + \frac{1+\epsilon}{1-\epsilon} (\frac{V}{10} + 1)$, which is less than $\frac{V}{4}$ for sufficiently large V .

To summarize, we have shown that the optimal contest for the max objective satisfies $v_1 \geq \frac{9V}{10} - 1$ whereas the optimal contest for the sum objective satisfies $v_1 < \frac{9V}{10} - 1$, so no single contest is optimal for both objectives.

5 Discussion

In this paper, we explored *behavioral* design—the idea that ‘non-standard’ agent behavior might have consequences for the optimal design of economic environments—via the specific problem of contest design for *simple* agents, who strategize only about whether or not to participate in a system rather than also strategizing about the input they provide to it.

Our results have potential practical implications for the design of contests as well as contest-like environments (such as online communities that hand out virtual rewards to top contributors) where the *quality* of a participant’s contribution is not the result of strategic deliberation. The structure of optimal contests in homogeneous cost environments suggests, for example, that a website might create stronger incentives for optimal contribution if it awards a (suitably-chosen) number of identical badges to top contributors instead of rewarding top contributors via recognition on a leaderboard, which might yield unequal rewards at different leaderboard ranks. In settings with heterogeneous costs, as in online crowdsourcing contests where participants may differ both in expertise and in cost (of effort or time) to contribution, our results show that a detail-free winner-take-all mechanism, while not necessarily optimal, is always approximately optimal—a guarantee which does not depend, for example, on the specific form of the cost-quality correlation, or the size of the contestant population, or the available total reward.

The differing structure of equilibria in the homogeneous and heterogeneous cost settings, as well as the example in §4.3 showing that the same contest cannot simultaneously (exactly) optimize the max and sum objectives with heterogeneous costs, suggest another interesting practical implication. The designer of an online community in which contribution is perceived as essentially equally costly by all agents does not have to choose whether her objective is to maximize participation (‘build a community’), or maximize the quality of the best contribution or the best few contributions (‘create excellent content’), or any combination thereof—as we saw in §3, all (increasing) objectives of the designer are maximized by the same reward structure in this setting. In contrast, the designer of a platform where users may have differing costs to participation—such as busy experts with a high time-cost to contribution versus users with lower expertise and lower costs—might have to make a choice about which objective she would like her platform to optimize, and design her choice of rewards accordingly.

5.1 Further directions

Our results answer the question of how to design optimal contests for simple agents—exactly, under homogeneous costs, and approximately, for heterogeneous costs, and suggest that ‘non-standard’ agent behavior can indeed have consequences for optimal design, as well as lead to interesting new theoretical questions. At the same time, our work also raises more questions than it answers—at various levels—and points to several avenues for further work.

Robustness. An intriguing family of open questions, in the specific setting of optimal contest design for simple agents, centers around the theme of robustness or simultaneous approximate optimality.

- **Robustness to *objective*:** We saw that the winner-take-all mechanism, which awards the entire prize of V to the top-ranked agent, is a 3-approximation for the max objective in the setting with arbitrary cost-quality correlations. Are there simple, detail-free, contests that are also simultaneously approximately optimal for other objectives that the designer might have, such as the sum of qualities?
- **Robustness to *parameters*:** Consider the setting with homogeneous costs, where we showed that optimal contests are *simple*, *i.e.*, award equal prizes to the k top-ranked agents for some k . A principal

without complete knowledge of the population parameters might choose the wrong number of prizes. How suboptimal can the ‘wrong simple contest’ be, and is there a contest which is nearly optimal for all population parameters¹⁷?

- **Robustness to *population*:** The model in our paper assumes that all agents are simple, whereas the past literature on contest design involved models where all agents make strategic effort choices. What if the population consists of a *heterogeneous* mix of agent types—those who make participation-only choices, and those who make strategic effort choices as well? Are there mechanisms that are robust to such heterogeneity in the *sets of alternatives* from which agents choose actions?

Exploring the robustness and suboptimality of simple contests is a practically relevant, although theoretically challenging, direction for further work, since a principal designing a real-world contest is more likely than not to suffer from uncertainty regarding population parameters, as well as have agents of differing types in her population.

Tournament design. A second question concerns optimal *tournament* design for simple agents. In this paper, we restricted ourselves to choosing amongst rank-order mechanisms where an agent’s reward depends only on her rank amongst all competitors. In general, of course, competitions need not be restricted to such reward structures: indeed, many real-life contests do split agents into subdivisions or categories, and award category-specific as well as overall prizes. Analogous to the reward allocation problems analyzed in [22] (which build on the techniques developed for optimal contest design in [21]) for agents who make strategic effort choices, what are optimal *tournament* designs for simple agents?

More generally, the notion of simple agents—agents who reason strategically about *whether*, but not about *how*, to participate in a system—seems to be a natural one in many scenarios. In addition to our running example of online content contribution platforms, settings such as crowdfunding websites (where investors might decide whether or not to contribute a recommended sum, rather than strategize over how much to contribute), online ratings, or voting (on content, or in participatory democracies) might elicit such simple behavior from at least some of the user base. How does the design of optimal mechanisms change when agents make participation-only choices?

While one might expect that a model with simpler agents ought to result in simpler mechanism design questions, we have seen in §4 that restricting agents to make discrete (binary) choices rather than selecting a level of effort from a continuous interval leads to qualitatively different equilibria and requires qualitatively different analysis techniques. This suggests that designing for simple agents—a problem that might arise in a variety of economic settings—may require the development of new theoretical tools and techniques.

Finally, a population of simple agents is one particular instance of a model for behavioral mechanism design, aimed at more effectively incentivizing desired outcomes in an environment by using more accurate models of agent behavior in that environment. A range of other deviations from idealized, or standard, models of behavior have also been discussed in the literature, ranging from workers’ preferences in online crowdsourcing labor markets [14], to a vast variety of biases and fallacies documented in the growing behavioral economics literature. The problems of developing better models of agent behavior in economic environments—an empirical and experimental question—and that of optimally designing mechanisms or markets for the real populations they are targeted at—a theoretical question—are closely intertwined, and constitute a promising and exciting direction for future research.

¹⁷Note that the answer to this last question might depend on the objective that the designer wants to optimize.

Acknowledgments

Robert Kleinberg acknowledges support from NSF award AF-0910940, a Microsoft Research New Faculty Fellowship, and a Google Research Grant.

References

- [1] Archak, N. and Sundarajan, A. (2009). Optimal design of crowdsourcing contests. In *International Conference on Information Systems*.
- [2] Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton Univ. Press.
- [3] Cavallo, R. and Jain, S. (2012). Efficient crowdsourcing contests. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [4] Chawla, S., Hartline, J., and Sivan, B. (2012). Optimal crowdsourcing contests. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [5] Che, Y.-K. and Gale, I. (2003). Optimal design of research tournaments. *American Economic Review*, 93(3):646–671.
- [6] DiPalantino, D. and Vojnović, M. (2009). Crowdsourcing and all-pay auctions. In *10th ACM conference on Electronic commerce (EC)*.
- [7] Easley, D. and Ghosh, A. (2013). Incentives, gamification, and game theory: An economic approach to badge design. In *Proc. 14th ACM Conf. on Electronic Commerce*, pages 359–376.
- [8] Ghosh, A. and Hummel, P. (2012). Implementing optimal outcomes in social computing. In *21st International World Wide Web Conference (WWW’12)*.
- [9] Ghosh, A. and Ligett, K. (2013). Privacy and coordination: Computing on databases with endogenous participation. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC ’13*.
- [10] Ghosh, A. and McAfee, R. (2012). Crowdsourcing with endogenous entry. In *21st International World Wide Web Conference (WWW’12)*.
- [11] Glazer, A. and Hassin, R. (1988). Optimal contests. *Economic Inquiry*, 26(1):133–143.
- [12] Glicksberg, I. L. (1952). A further generalization of the kakutani fixed point theorem, with applications to nash equilibrium points. *Proc. Amer. Math. Soc.*, 3(1):170–174.
- [13] Green, J. R. and Stokey, N. L. (1983). A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–364.
- [14] Horton, J. and Chilton, L. (2010). The labor economics of paid crowdsourcing. In *Proceedings of the 11th ACM Conference on Electronic Commerce, EC ’10*.
- [15] Immorlica, N., Stoddard, N., and Syrgkanis, V. (2013). Social status and badge design. eprint arXiv:1312.2299.

- [16] Jeon, G. Y., Kim, Y.-M., and Chen, Y. (2010). Re-examining price as a predictor of answer quality in an online Q&A site. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10.
- [17] Kraut, R. and Resnick, P. (2012). *Building successful online communities: Evidence-based social design*. Cambridge University Press.
- [18] Larson, K. and Sandholm, T. (2005). Mechanism design and deliberative agents. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [19] Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864.
- [20] Milgrom, P. R. and Weber, R. J. (1985). Distributional strategies for games with incomplete information. *Math. of O. R.*, 10(4):619–632.
- [21] Moldovanu, B. and Sela, A. (2001). Increasing effort through softening incentives in contests. *American Economic Review*, 91(3):542–558.
- [22] Moldovanu, B. and Sela, A. (2006). Contest architecture. *J. Econ. Theory*, 126(1):70–97.
- [23] Moldovanu, B., Sela, A., and Shi, X. (2007). Contests for status. *Journal of Political Economy*, 115(2):338–363.
- [24] Motwani, R. and Raghavan, P. (1995). *Randomized Algorithms*. Cambridge Univ. Press.
- [25] Samuel-Cahn, E. (1984). Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12(4):1213–1216.
- [26] Taylor, C. R. (1995). Digging for golden carrots: An analysis of research tournaments. *American Economic Review*, 85(4):872–890.
- [27] Wright, J. and Leyton-Brown, K. (2014). Level-0 meta-models for predicting human behavior in games. In *Proc. 15th ACM Conference on Economics and Computation (EC'14)*.

A Proofs from Section 3

This section provides proofs deferred from §3. We begin by restating and proving Proposition 3.1.

Proposition 3.1. *For any monotone rank-order mechanism \mathcal{M} , there is a symmetric equilibrium in threshold strategies where every agent with quality $q_i \geq q^*$ participates, and every agent with quality $q_i < q^*$ does not participate; the equilibrium threshold q^* satisfies $c_{\mathcal{M}}(1 - F(q^*)) = c$. This is the unique symmetric equilibrium of \mathcal{M} .*

Proof. Let π be any symmetric pure-strategy equilibrium of \mathcal{M} and let $q_0 = \inf\{q \mid \pi(q, c) = 1\}$. By the definition of q_0 , for any $q > q_0$ there is some $q_1 \leq q$ such that $\pi(q_1, c) = 1$. Since π is an equilibrium, this implies that an agent with quality q_1 earns an expected prize greater than or equal to c in a contest with $n - 1$ potential opponents each playing strategy π . Since \mathcal{M} is monotone, the expected prize given to an agent is a non-decreasing function of her output quality, so that an agent with quality $q > q_1$ earns an expected prize

greater than or equal to her cost: since ties are broken in favor of participation, this means that $\pi(q, c) = 1$; and this holds for all $q > q_0$. By the definition of q_0 , $\pi(q, c) = 0$ for all $q < q_0$; finally, our rule that ties are broken in favor of participation, combined with our assumption that the distribution of q has no point-mass at q_0 , ensures that $\pi(q_0, c) = 1$. Hence, π is a threshold-strategy equilibrium with threshold $q^* = q_0$.

Next we characterize the equilibrium threshold q^* , and argue that it is the solution to $c_{\mathcal{M}}(1 - F(q^*)) = c$. Note that when all agents play strategy π , $1 - F(q^*)$ is the probability that a given agent participates and contributes an output with quality greater than or equal to q^* . Recalling our definition of $c_{\mathcal{M}}$, and noting that when agents play according to π they *never* contribute outputs with quality less than q^* , we can now conclude that $c_{\mathcal{M}}(1 - F(q^*))$ is equal to the expected prize earned by a contestant with quality $q \leq q^*$ when the opponents play according to π . If this were less than c , then it could not be the case that $\pi(q^*, c) = 1$ in equilibrium. On the other hand, if $c_{\mathcal{M}}(1 - F(q^*))$ were greater than c , then the expected prize earned by any participating agent with quality $q \leq q^*$ is greater than c and consequently all such agents should participate in equilibrium, contradicting the fact that $\pi(q, c) = 0$ for $q < q^*$. Since each of the inequalities $c_{\mathcal{M}}(1 - F(q^*)) < c$ and $c_{\mathcal{M}}(1 - F(q^*)) > c$ leads to a contradiction, it must be the case that $c_{\mathcal{M}}(1 - F(q^*)) = c$. \square

Continuing with deferred proofs from §3, let us restate and prove Proposition 3.2.

Proposition 3.2. *Fix a total prize, V , and number of potential contestants, n . Let \mathcal{M}^j denote a simple contest that awards equal prizes V/j to the top j ranks.*

1. *There is a decreasing sequence $V = c_1 > c_2 > \dots > c_n > c_{n+1} = 0$, such that for any participation cost $c < V$, the simple contest \mathcal{M}^j is optimal if and only if $c_{j+1} \leq c \leq c_j$.*
2. *For $1 < j \leq n$, the equation $c_{\mathcal{M}^j}(p) = c_{\mathcal{M}^{j-1}}(p)$ has a unique solution p_j in the interval $(0, 1)$, and $c_j = c_{\mathcal{M}^j}(p_j) = c_{\mathcal{M}^{j-1}}(p_j)$.*

Proof. Let us compare the simple contests \mathcal{M}^j and \mathcal{M}^{j-1} , for $1 < j \leq n$. For a given participation cost, c , let q^j and q^{j-1} denote the equilibrium thresholds of \mathcal{M}^j and \mathcal{M}^{j-1} , respectively. The probability that an individual agent participates in the equilibrium of \mathcal{M}^j (resp. \mathcal{M}^{j-1}) will be denoted by $p^j = 1 - F(q^j)$, resp. $p^{j-1} = 1 - F(q^{j-1})$. From Proposition 3.1 we know that $c = c_{\mathcal{M}^j}(p^j) = c_{\mathcal{M}^{j-1}}(p^{j-1})$, which implies

$$0 = c_{\mathcal{M}^j}(p^j) - c_{\mathcal{M}^{j-1}}(p^{j-1}) = [c_{\mathcal{M}^j}(p^j) - c_{\mathcal{M}^{j-1}}(p^j)] + [c_{\mathcal{M}^{j-1}}(p^j) - c_{\mathcal{M}^{j-1}}(p^{j-1})]. \quad (14)$$

Combining this equation with the fact that the function $c_{\mathcal{M}^{j-1}}$ is strictly decreasing (Lemma 3.1) we obtain the following chain of equivalences:

$$p^j > p^{j-1} \iff c_{\mathcal{M}^{j-1}}(p^j) < c_{\mathcal{M}^{j-1}}(p^{j-1}) \iff c_{\mathcal{M}^j}(p^j) - c_{\mathcal{M}^{j-1}}(p^j) > 0. \quad (15)$$

Thus, the question of which mechanism, \mathcal{M}^j or \mathcal{M}^{j-1} , has a higher rate of participation in equilibrium reduces to evaluating whether $c_{\mathcal{M}^j} - c_{\mathcal{M}^{j-1}}$ is positive or negative at p^j . To understand the sign of the

function $c_{\mathcal{M}^j} - c_{\mathcal{M}^{j-1}}$ on the interval $[0, 1]$, the following calculation is crucial.

$$\begin{aligned}
c_{\mathcal{M}^j(p)} - c_{\mathcal{M}^{j-1}(p)} &= \frac{V}{j} \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} - \frac{V}{j-1} \sum_{k=1}^{j-1} \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \\
&= \frac{V}{j} \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j} + \sum_{k=1}^{j-1} \left(\frac{V}{j} - \frac{V}{j-1} \right) \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \\
&= \frac{V}{j} (1-p)^{n-1} \binom{n-1}{j-1} \left(\frac{p}{1-p} \right)^{j-1} - \frac{V}{j(j-1)} (1-p)^{n-1} \sum_{k=1}^{j-1} \binom{n-1}{k-1} \left(\frac{p}{1-p} \right)^{k-1} \\
&= \frac{V}{j(j-1)} (1-p)^{n-1} \left[(j-1) \binom{n-1}{j-1} \left(\frac{p}{1-p} \right)^{j-1} - \sum_{k=1}^{j-1} \binom{n-1}{k-1} \left(\frac{p}{1-p} \right)^{k-1} \right]
\end{aligned}$$

Letting $x = \frac{p}{1-p}$ we have derived

$$c_{\mathcal{M}^j(p)} - c_{\mathcal{M}^{j-1}(p)} = \frac{V}{j(j-1)} (1-p)^{n-1} \left[(j-1) \binom{n-1}{j-1} x^{j-1} - \sum_{k=1}^{j-1} \binom{n-1}{k-1} x^{k-1} \right] \quad (16)$$

The polynomial $(j-1) \binom{n-1}{j-1} x^{j-1} - \binom{n-1}{j-2} x^{j-2} - \binom{n-1}{j-3} x^{j-3} - \dots - \binom{n-1}{1} x - 1$ appearing on the right side of (16) is negative at $x = 0$, tends to $+\infty$ as $x \rightarrow \infty$, and by Descartes' Rule of Signs, it has a single positive real root, r . Accordingly, there is a unique $p_j = \frac{r}{r+1}$ in the interval $(0, 1)$ satisfying the equation $c_{\mathcal{M}^j}(p_j) = c_{\mathcal{M}^{j-1}}(p_j)$, and the quantity $c_{\mathcal{M}^j}(p) - c_{\mathcal{M}^{j-1}}(p)$ is positive if and only if $p > p_j$. As we vary the participation cost, c , the equilibrium participation probability of the simple contest \mathcal{M}^j , namely p^j , is a decreasing function of c . Therefore, letting $c_j = c_{\mathcal{M}^j}(p_j) = c_{\mathcal{M}^{j-1}}(p_j)$, we have the following chain of equivalences.

$$p^j > p^{j-1} \iff c_{\mathcal{M}^j}(p^j) - c_{\mathcal{M}^{j-1}}(p^j) > 0 \iff p^j > p_j \iff c < c_j. \quad (17)$$

The proposition claims that the numbers (c_j) form a decreasing sequence. We now turn to the proof of this fact. Recalling that the equation $c_j = c_{\mathcal{M}^j}(p_j) = c_{\mathcal{M}^{j-1}}(p_j)$ holds for all $j \geq 2$, we find that

$$c_j - c_{j+1} = c_{\mathcal{M}^j}(p_j) - c_{\mathcal{M}^j}(p_{j+1}). \quad (18)$$

The function $x(p) = \frac{p}{1-p}$ is an increasing function of p , and $c_{\mathcal{M}^j}(p)$ is a decreasing function of p . This justifies the second in the following pair of equivalences.

$$c_j > c_{j+1} \stackrel{(18)}{\iff} c_{\mathcal{M}^j}(p_j) > c_{\mathcal{M}^j}(p_{j+1}) \iff x(p_j) < x(p_{j+1}), \quad (19)$$

We are left with proving that the RHS of (19) is valid. For every $j \geq 2$ define the polynomial

$$Q_j(x) = (j-1) \binom{n-1}{j-1} x^{j-1} - \sum_{k=1}^{j-1} \binom{n-1}{k-1} x^{k-1}. \quad (20)$$

We know from (16) and the ensuing paragraph that the sign of $Q_j(x)$ is the same as the sign of $x - x(p_j)$, so to prove that $x(p_j) < x(p_{j+1})$, it suffices to show that $Q_j(x(p_{j+1})) > 0$. The easy calculation

$$Q_{j+1}(x) - Q_j(x) = j \binom{n-1}{j} x^j - j \binom{n-1}{j-1} x^{j-1} = j \binom{n-1}{j} x^{j-1} \left(x - \frac{j}{n-j} \right), \quad (21)$$

combined with the fact that $Q_{j+1}(x(p_{j+1})) = 0$, reveals that $Q_j(x(p_{j+1})) > 0$ if and only if $x(p_{j+1}) < \frac{j}{n-j}$. Recalling that $x(p_{j+1}) < x$ if and only if $Q_{j+1}(x) > 0$, we are left with showing that $Q_{j+1}\left(\frac{j}{n-j}\right) > 0$. The calculation is as follows.

$$\begin{aligned}
Q_{j+1}\left(\frac{j}{n-j}\right) &= j \binom{n-1}{j} \left(\frac{j}{n-j}\right)^j - \sum_{k=1}^j \binom{n-1}{k-1} \left(\frac{j}{n-j}\right)^{k-1} \\
&= \sum_{k=1}^j \left[\binom{n-1}{j} \left(\frac{j}{n-j}\right)^j - \binom{n-1}{k-1} \left(\frac{j}{n-j}\right)^{k-1} \right] \\
&= \binom{n-1}{j} \left(\frac{j}{n-j}\right)^j \sum_{k=1}^j \left[1 - \frac{(n-1)!}{(k-1)!(n-k)!} \cdot \frac{j!(n-1-j)!}{(n-1)!} \cdot \left(\frac{j}{n-j}\right)^{k-1-j} \right] \\
&= \binom{n-1}{j} \left(\frac{j}{n-j}\right)^j \sum_{k=1}^j \left[1 - \frac{j!}{(k-1)!} \cdot \frac{(n-1-j)!}{(n-k)!} \cdot \left(\frac{j}{n-j}\right)^{k-1-j} \right] \\
&= \binom{n-1}{j} \left(\frac{j}{n-j}\right)^j \sum_{k=1}^j \left[1 - \left(\prod_{i=k}^j \frac{i}{n-i} \cdot \frac{n-j}{j} \right) \right] \\
&= \binom{n-1}{j} \left(\frac{j}{n-j}\right)^j \sum_{k=1}^j \left[1 - \prod_{i=k}^j \frac{in-ij}{jn-ij} \right] \\
&= \binom{n-1}{j} \left(\frac{j}{n-j}\right)^j \sum_{k=1}^{j-1} \left[1 - \prod_{i=k}^j \frac{in-ij}{jn-ij} \right]
\end{aligned}$$

Recall that we are assuming $j > 1$, so the sum in the final line is over a non-empty set of indices k . Each term in the product on the final line is less than or equal to 1, at least one of the terms is strictly less than 1, so the product is strictly less than 1. This means that each summand on the final line is strictly positive, so the entire quantity on the final line is strictly positive, as desired. This completes the proof that $Q_{j+1}\left(\frac{j}{n-j}\right) > 0$, and as we have seen this implies that $c_j > c_{j+1}$ for $j > 1$. To complete the proof that the entire sequence c_1, c_2, \dots, c_{n+1} is strictly decreasing, we need only prove that $c_1 > c_2$. This inequality follows easily, once we recall that $c_1 = V$ while

$$c_2 = c_{\mathcal{M}^2}(p_2) = \frac{V}{2} \sum_{k=1}^2 \binom{n-1}{k-1} p_2^{k-1} (1-p_2)^{n-k} \leq \frac{V}{2} \sum_{k=1}^n \binom{n-1}{k-1} p_2^{k-1} (1-p_2)^{n-k} = \frac{V}{2}.$$

It remains to show that \mathcal{M}^j is an optimal simple contest if and only if $c_{j+1} \leq c \leq c_j$. Recall that for every $j > 2$, \mathcal{M}^j has a higher rate of participation than \mathcal{M}^{j-1} , in equilibrium, if and only if $c < c_j$. This immediately shows that $c_{j+1} \leq c \leq c_j$ is a necessary condition for \mathcal{M}^j to be optimal, as otherwise either \mathcal{M}^{j+1} or \mathcal{M}^{j-1} is superior to \mathcal{M}^j . To prove that it is also a sufficient condition, we will prove that \mathcal{M}^j is strictly optimal whenever $c_{j+1} < c < c_j$; an easy continuity argument then confirms that \mathcal{M}^j is weakly optimal whenever $c_{j+1} \leq c \leq c_j$.

Suppose that $c_{j+1} < c < c_j$ and for every $k = 1, \dots, n$ let p^k denote the equilibrium participation probability for \mathcal{M}^k . For $2 \leq k \leq j$ we have the string of inequalities $c \leq c_j < c_{j-1} < \dots < c_{k+1} < c_k$, which means that $c < c_k$ and hence that $p^{k-1} < p^k$. Stringing these inequalities together for all k in the range $2, \dots, j$, we find that $p^1 < p^2 < \dots < p^{j-1} < p^j$, i.e., \mathcal{M}^j has a higher rate of participation, in

equilibrium, than \mathcal{M}^k for any $k < j$. The argument for $k > j$ is the same with the inequality signs reversed. For $j < k \leq n$ we have the string of inequalities $c \geq c_{j+1} > c_{j+2} \cdots > c_{k-1} > c_k$ which means that $c > c_k$ and hence that $p^{k-1} > p^k$. Stringing these inequalities together for all k in the range $j+1, \dots, n$, we find that $p^j > p^{j+1} > \dots > p^{n-1} > p^n$, i.e., \mathcal{M}^j has a higher rate of participation, in equilibrium, than \mathcal{M}^k for any $k > j$. \square

Proposition 3.3 makes three assertions about optimal contests and their equilibria in the setting of homogeneous costs. We now restate and prove the first two parts of the proposition. The following subsection presents the proof of the third part.

Proposition 3.3 (parts 1-2). *Fix V, c , and let $j^{*,n}$ denote the number of prizes offered by the optimal contest with population size n . Let $p^{*,n}$ denote the probability that an individual agent participates in the equilibrium of that contest, and let $\lambda^{*,n} = n \cdot p^{*,n}$ denote the expected number of participants in that equilibrium.*

1. *The winner-take-all contest is optimal if and only if $\frac{V}{c} \leq \left(1 + \frac{1}{n-1}\right)^{n-1}$. In particular, for $\frac{V}{c} > e$ a winner-take-all contest is never optimal, whereas for $\frac{V}{c} < e$ it is optimal whenever the population size, n , is sufficiently large.*
2. *As $n \rightarrow \infty$, the sequence $\lambda^{*,n}$ converges to the unique $\lambda^* > 0$ satisfying*

$$\max_{1 \leq j \leq V/c} \left\{ \frac{V}{j} \sum_{k=0}^{j-1} \frac{e^{-\lambda} \lambda^k}{k!} \right\} = c. \quad (22)$$

Excluding a countable set of values for c , the maximum in (7) is achieved at a unique value j^ , and the sequence $j^{*,n}$ converges to j^* as $n \rightarrow \infty$.*

Proof. Proposition 3.2 tells us that the winner-take-all contest is optimal if and only if $c \geq c_2 = c_{\mathcal{M}^1}(p)$, where p is the unique solution of the equation $c_{\mathcal{M}^1}(p) = c_{\mathcal{M}^2}(p)$. We now calculate:

$$\begin{aligned} c_{\mathcal{M}^1}(p) - c_{\mathcal{M}^2}(p) &= V(1-p)^{n-1} - \frac{V}{2} [(1-p)^{n-1} + (n-1)p(1-p)^{n-2}] \\ &= \frac{V}{2}(1-p)^{n-2} [1 - np]. \end{aligned}$$

Hence $p = \frac{1}{n}$ is the unique solution to $c_{\mathcal{M}^1}(p) = c_{\mathcal{M}^2}(p)$. Now we find that

$$c_2 = c_{\mathcal{M}^1}(p) = V(1-p)^{n-1} = V \left(1 - \frac{1}{n}\right)^{n-1} \quad (23)$$

so the requirement $c \geq c_2$ can be rewritten as $c \geq V \left(\frac{n-1}{n}\right)^{n-1}$, which in turn can be rewritten as $\frac{V}{c} \leq \left(1 + \frac{1}{n-1}\right)^{n-1}$, matching the criterion for optimality of winner-take-all claimed in part 1 of the proposition.

The observation that $\left(1 + \frac{1}{n-1}\right)^{n-1}$ is always less than e , and converges to e as $n \rightarrow \infty$, supplies the justification for the last sentence of part 1.

Moving on to the second part of the proposition, let $p^{j,n}$ denote the equilibrium participation probability in the simple contest \mathcal{M}^j when the population size is n . If $j > \nu = V/c$ then $V/j < c$, i.e., winning a prize in \mathcal{M}^j is not worth enough to cover the cost of participating. Hence, $p^{j,n} = 0$ if $j > \nu$. Using the fact that optimal contests are simple and that a simple contest with $j > \nu$ prizes has no participation, we find that

$$c = c^*(p^{*,n}) = \max_{1 \leq j \leq \nu} c_{\mathcal{M}^j}^n(p^{*,n}), \quad (24)$$

where $c_{\mathcal{M}^j}^n(\cdot)$ denotes the function $c_{\mathcal{M}}(\cdot)$ defined in Equation (1), specialized to population size n and contest $\mathcal{M} = \mathcal{M}^j$. Under the change of variables $\lambda = \lambda(p) = n \cdot p$, the function $c_{\mathcal{M}^j}^n(p)$ is expressed by the formula

$$c_{\mathcal{M}^j}^n(p) = c_{\mathcal{M}^j}^n\left(\frac{\lambda}{n}\right) = \frac{V}{j} \sum_{k=0}^{j-1} \binom{n-1}{k} \left(\frac{\lambda}{n}\right)^k \left(1 - \frac{\lambda}{n}\right)^{n-1-k}. \quad (25)$$

For any fixed j , as $n \rightarrow \infty$ the RHS of (25) converges uniformly as a function of $\lambda \in [0, \nu]$. Indeed, for all fixed $k \in \mathbb{N}$,

$$\lim_{n \rightarrow \infty} \left(1 - \frac{\lambda}{n}\right)^{n-1-k} = \left[\lim_{n \rightarrow \infty} \left(1 - \frac{\lambda}{n}\right)^n\right] \cdot \left[\lim_{n \rightarrow \infty} \left(1 - \frac{\lambda}{n}\right)^{-1-k}\right] = e^{-\lambda} \quad (26)$$

$$\lim_{n \rightarrow \infty} \binom{n-1}{k} \left(\frac{\lambda}{n}\right)^k = \lim_{n \rightarrow \infty} \prod_{i=1}^k \left(\frac{n-i}{i} \cdot \frac{\lambda}{n}\right) = \frac{\lambda^k}{k!} \left(\lim_{n \rightarrow \infty} \prod_{i=1}^k \frac{n-i}{n}\right) = \frac{\lambda^k}{k!} \quad (27)$$

$$\lim_{n \rightarrow \infty} c_{\mathcal{M}^j}^n\left(\frac{\lambda}{n}\right) = \frac{V}{j} \sum_{k=0}^{j-1} \frac{e^{-\lambda} \lambda^k}{k!}, \quad (28)$$

and the convergence on each line of (26)-(28) is uniform over $\lambda \in [0, \nu]$.

Let c_j^∞ denote the function appearing on the right side of (28), i.e. $c_j^\infty(\lambda) = \frac{V}{j} \sum_{k=0}^{j-1} \frac{e^{-\lambda} \lambda^k}{k!}$. This is a continuous and strictly decreasing function of λ , for all $\lambda \geq 0$. To verify this, we can simply take the derivative with respect to λ .

$$\begin{aligned} \frac{d}{d\lambda} \left[\sum_{k=0}^{j-1} \frac{e^{-\lambda} \lambda^k}{k!} \right] &= -e^{-\lambda} + \sum_{k=1}^{j-1} \left[\frac{e^{-\lambda} \lambda^{k-1}}{(k-1)!} - \frac{e^{-\lambda} \lambda^k}{k!} \right] \\ &= e^{-\lambda} \left[\sum_{i=0}^{j-2} \frac{\lambda^i}{i!} - \sum_{k=0}^{j-1} \frac{\lambda^k}{k!} \right] = -\frac{e^{-\lambda} \lambda^{j-1}}{(j-1)!}. \end{aligned} \quad (29)$$

It is easy to see that $c_j^\infty(0) = \frac{V}{j}$ and that $\lim_{\lambda \rightarrow \infty} c_j^\infty(\lambda) = 0$. Combined with the fact that c_j^∞ is continuous and strictly decreasing, this means it has a continuous and strictly decreasing inverse function $\lambda^j(c)$, defined for $0 < c \leq \frac{V}{j}$.

Let $\lambda^{j,n} = n \cdot p^{j,n}$. If $n(1) < n(2) < \dots$ is any infinite sequence of indices such that $\lim_{i \rightarrow \infty} \lambda^{j,n(i)}$ exists, then the uniform convergence in Equation (28) combined with the equation $c_{\mathcal{M}^j}^n(p^{j,n}) = c$ which is valid for all j, n , imply that

$$c = \lim_{i \rightarrow \infty} c_{\mathcal{M}^j}^{n(i)}(p^{j,n(i)}) = \lim_{i \rightarrow \infty} c_{\mathcal{M}^j}^{n(i)}\left(\frac{\lambda^{j,n(i)}}{n(i)}\right) = c_j^\infty\left(\lim_{i \rightarrow \infty} \lambda^{j,n(i)}\right). \quad (30)$$

Recalling that the inverse function of c_j^∞ is denoted by λ^j , this means that $\lim_{i \rightarrow \infty} \lambda^{j,n(i)} = \lambda^j(c)$.

Recall now that $\lambda^{j,n}$ can be interpreted as the expected number of participants in the equilibrium of contest \mathcal{M}^j when the population size is n . This means that the expected total participation cost is $\lambda^{j,n} \cdot c$. Individual rationality implies that $\lambda^{j,n} \cdot c \leq V$, hence $\lambda^{j,n} \leq V/c = \nu$. Thus, every element of the infinite sequence $(\lambda^{j,n})_{n=1}^\infty$ belongs to the interval $[0, \nu]$. Earlier we derived that for every infinite subsequence that has a limit point, this limit point must be equal to $\lambda^j(c)$. It follows that the entire sequence $\lambda^{j,n}$ must converge to $\lambda^j(c)$, as otherwise there would be some $\varepsilon > 0$ and an infinite subsequence belonging to the set $[0, \lambda^j(c) - \varepsilon] \cup [\lambda^j(c) + \varepsilon, \nu]$. That set is compact, so the infinite subsequence would have a limit point in the set, contradicting our earlier derivation that every subsequential limit point must be equal to $\lambda^j(c)$.

Earlier we noted that $p^{j,n} = 0$ for $j > \nu$; hence $p^{*,n} = \max_{1 \leq j \leq \nu} \{p^{j,n}\}$ and $\lambda^{*,n} = \max_{1 \leq j \leq \nu} \{\lambda^{j,n}\}$. This implies

$$\lim_{n \rightarrow \infty} \lambda^{*,n} = \lim_{n \rightarrow \infty} \left(\max_{1 \leq j \leq \nu} \{\lambda^{j,n}\} \right) = \max_{1 \leq j \leq \nu} \left\{ \lim_{n \rightarrow \infty} \lambda^{j,n} \right\} = \max_{1 \leq j \leq \nu} \{\lambda^j(c)\}. \quad (31)$$

We must show that the right side of (33) is equal to the unique $\lambda^* > 0$ satisfying Equation (22). Using notation we have introduced in this proof, that equation can be rewritten as

$$\max_{1 \leq j \leq \nu} \{c_j^\infty(\lambda^*)\} = c. \quad (32)$$

The function $h(\lambda) = \max_{1 \leq j \leq \nu} \{c_j^\infty(\lambda)\}$, being a pointwise maximum of continuous and strictly decreasing functions, is also continuous and strictly decreasing. Thus, there is a unique $\lambda^* = \lambda^*(c)$ satisfying (32). Let j^* be the value of j that achieves the maximum in the formula defining $h(\lambda^*)$, so that $c = h(\lambda^*) = c_{j^*}^\infty(\lambda^*)$. Applying the function λ^{j^*} to both sides of this equation yields $\lambda^{j^*}(c) = \lambda^*$, and therefore $\lambda^* \leq \max_{1 \leq j \leq \nu} \lambda^j(c)$. The inequality cannot be strict, because if $\lambda^* < \lambda^j(c)$ for some j , then we can apply the strictly decreasing function c_j^∞ to both sides and conclude that $c_{j^*}^\infty(\lambda^*) > c_j(\lambda^j(c)) = c$, contradicting Equation (32).

The proposition also asserts that with the exception of countably many values of c , the maximum in (22) is achieved by a unique value of j . To show this, consider for any $j \neq \ell$ the equation

$$\frac{V}{j} \sum_{k=0}^{j-1} \frac{e^{-\lambda} \lambda^k}{k!} - \frac{V}{\ell} \sum_{k=0}^{\ell-1} \frac{e^{-\lambda} \lambda^k}{k!} = 0, \quad (33)$$

which must be satisfied by $\lambda^*(c)$ whenever the maximum in (22) is achieved at both j and ℓ . For any two solutions λ_0, λ_1 of (33), there must be a λ in the open interval (λ_0, λ_1) such that the derivative of the LHS of (33) equals zero at λ . We know, from Equation (29), that the derivative of the LHS of (33) is equal to $V e^{-\lambda} \left(\frac{\lambda^{\ell-1}}{\ell!} - \frac{\lambda^{j-1}}{j!} \right)$, which equates to zero only at $\lambda = (\ell!/j!)^{1/(\ell-j)}$. Thus, Equation (33) can have at most two solutions for any fixed j, ℓ . Taking the union over all pairs j, ℓ , we find that there are at most countably many solutions to (33). For any c such that $\lambda^*(c)$ does not belong to this countable set, the maximum in (22) is attained at a unique $j^* = j^*(c)$.

Finally we must prove that when c lies outside this countable set, the sequence $j^{*,n}$ converges to j^* as $n \rightarrow \infty$. Indeed, we have seen that each element of the sequence $(j^{*,n})_{n=1}^\infty$ is an integer in the range $[1, \nu]$, so there must be at least one $j \in [1, \nu]$ that occurs infinitely often in the sequence. We will show that any such j must equal j^* , from which it follows immediately that $j^{*,n} \rightarrow j^*$ as $n \rightarrow \infty$.

Now suppose $n(1) < n(2) < \dots$ is an infinite sequence such that $j^{*,n(i)} = j$ for all $i \geq 1$. We have

$$\begin{aligned} \forall i \geq 1 \quad c_{\mathcal{M}^j}^{n(i)} \left(\frac{\lambda^{*,n(i)}}{n} \right) &= c_{\mathcal{M}^j}^{n(i)} \left(\frac{\lambda^{j,n(i)}}{n} \right) = c_{\mathcal{M}^j}^{n(i)} (p^{j,n(i)}) = c \\ c_j^\infty(\lambda^*) &= c_j^\infty \left(\lim_{i \rightarrow \infty} \lambda^{*,n(i)} \right) = \lim_{i \rightarrow \infty} c_{\mathcal{M}^j}^{n(i)} \left(\frac{\lambda^{*,n(i)}}{n} \right) = c \end{aligned} \quad (34)$$

where the last line is justified by the uniform convergence guarantee in Equation (28). We know that $c_{j^*}^\infty(\lambda^*) = c$, and our assumption on c implies that for every $j' \neq j^*$, $c_{j'}^\infty(\lambda^*) < c$. Hence, Equation (34) implies that $j = j^*$, which completes the proof of the proposition. \square

A.1 Proof of Proposition 3.3, part 3

The proof of Proposition 3.3, part 3 depends on some rather sharp estimates of binomial coefficients. To derive these estimates, we begin by deriving a non-asymptotic version of Stirling's approximation for $n!$.

Lemma A.1. *For any positive integer n ,*

$$n \ln(n) - n + \frac{1}{2} \ln(n) + \frac{2}{3} < \ln(n!) < n \ln(n) - n + \frac{1}{2} \ln(n) + 1. \quad (35)$$

Proof. For all k and all $t \in (0, 1)$ we have

$$\ln(k) + t(\ln(k+1) - \ln(k)) < \ln(k+t) < \ln(k) + \frac{t}{k},$$

where the left inequality is derived from the strict concavity of the natural logarithm function, and the right inequality is derived from strict concavity along with the fact that the derivative of the natural logarithm at k is $\frac{1}{k}$. Integrating with respect to t and applying the substitution $x = k + t$, we find that

$$\ln(k) + \frac{1}{2}(\ln(k+1) - \ln(k)) < \int_k^{k+1} \ln(x) dx < \ln(k) + \frac{1}{2k}. \quad (36)$$

Now, summing over $k = 1, \dots, n-1$,

$$\ln(n!) - \frac{1}{2} \ln(n) < \int_1^n \ln(x) dx < \ln(n!) - \ln(n) + \frac{1}{2} \sum_{k=1}^{n-1} \frac{1}{k} < \ln(n!) - \frac{1}{2} \ln(n) + \frac{\gamma}{2}, \quad (37)$$

where $\gamma = 0.577\dots$ is the Euler-Mascheroni constant and we have used the fact that $\sum_{k=1}^{n-1} \frac{1}{k} < \ln(n) + \gamma$. Rearranging terms in (37) and using the fact that $\int_1^n \ln(x) dx = n \ln(n) - n + 1$, we derive

$$n \ln(n) - n + \frac{1}{2} \ln(n) + 1 - \frac{\gamma}{2} < \ln(n!) < n \ln(n) - n + \frac{1}{2} \ln(n) + 1 \quad (38)$$

The lemma follows by observing that $\frac{2}{3} < 1 - \frac{\gamma}{2}$. \square

Lemma A.2. *For any positive integers n and $j \leq \frac{n}{2}$ and any $p \in (0, 1)$,*

$$\binom{n}{j} p^j (1-p)^{n-j} > \frac{1}{4} \cdot \sqrt{\frac{1}{j}} \cdot \exp\left(-\frac{2(j-pn)^2}{pn}\right). \quad (39)$$

Furthermore if $j > pn$ then

$$\binom{n}{j} p^j (1-p)^{n-j} < \exp\left(-\frac{(j-pn)^2 - 2}{2j}\right). \quad (40)$$

Proof. Using Lemma A.1 and the formula $\binom{n}{j} = \frac{n!}{j!(n-j)!}$, we obtain the bound

$$\begin{aligned} \ln \binom{n}{j} &= \ln(n!) - \ln(j!) - \ln((n-j)!) \\ &> n \ln(n) - j \ln(j) - (n-j) \ln(n-j) + \frac{1}{2} \ln\left(\frac{n}{j(n-j)}\right) - \frac{4}{3} \\ &\geq j \ln\left(\frac{n}{j}\right) + (n-j) \ln\left(\frac{n}{n-j}\right) + \frac{1}{2} \ln\left(\frac{1}{j}\right) - \frac{4}{3}. \end{aligned} \quad (41)$$

Therefore the logarithm of the left side of (39) satisfies

$$\ln \binom{n}{j} + j \ln(p) + (n-j) \ln(1-p) > j \ln \left(\frac{pn}{j} \right) + (n-j) \ln \left(\frac{(1-p)n}{n-j} \right) + \frac{1}{2} \ln \left(\frac{1}{j} \right) - \frac{4}{3} \quad (42)$$

Focusing attention on the first two terms on the right side of (42), we may perform the following manipulation making use of the identity $\ln \left(\frac{1}{x} \right) = -\ln(x) \geq 1-x$, which holds for all $x > 0$.

$$\begin{aligned} j \ln \left(\frac{pn}{j} \right) + (n-j) \ln \left(\frac{(1-p)n}{n-j} \right) &= (j-pn) \ln \left(\frac{pn}{j} \right) + (pn) \ln \left(\frac{pn}{j} \right) + \\ &\quad (pn-j) \ln \left(\frac{(1-p)n}{n-j} \right) + (n-pn) \ln \left(\frac{(1-p)n}{n-j} \right) \\ &\geq (j-pn) \ln \left(\frac{pn(n-j)}{j(1-p)n} \right) + pn \left(1 - \frac{j}{pn} \right) + (n-pn) \left(1 - \frac{n-j}{n-pn} \right) \\ &= (j-pn) \ln \left(\frac{pn-pj}{j-pj} \right) + pn - j + n - pn - (n-j) \\ &\geq (j-pn) \left(1 - \frac{j-pj}{pn-pj} \right) = -\frac{(j-pn)^2}{p(n-j)}. \end{aligned} \quad (43)$$

Combining (43) with (42), we obtain

$$\begin{aligned} \ln \binom{n}{j} + j \ln(p) + (n-j) \ln(1-p) &> -\frac{(j-pn)^2}{p(n-j)} - \frac{1}{2} \ln(j) - \frac{4}{3} \\ &> -\frac{2(j-pn)^2}{pn} - \frac{1}{2} \ln(j) - \frac{4}{3}, \end{aligned} \quad (44)$$

where the second inequality follows because $n-j \geq \frac{n}{2}$. Exponentiating both sides of (44) and using the fact that $e^{4/3} < 4$, we obtain the inequality (39) stated in the lemma.

For the second part of the lemma, let $\ell = \lfloor pn \rfloor$. The expression $\binom{n}{\ell} p^\ell (1-p)^{n-\ell}$ is the probability of observing ℓ heads when tossing a bias- p coin n times. Using the trivial fact that this probability is at most 1, we obtain the bound

$$\begin{aligned} \binom{n}{j} p^j (1-p)^{n-j} &\leq \frac{\binom{n}{j} p^j (1-p)^{n-j}}{\binom{n}{\ell} p^\ell (1-p)^{n-\ell}} = \frac{\ell! (n-\ell)! p^j (1-p)^{n-j}}{j! (n-j)! p^\ell (1-p)^{n-\ell}} \\ &= \prod_{k=\ell+1}^j \frac{(n-k)p}{k(1-p)} = \prod_{k=\ell+1}^j \left(1 - \frac{k-pn}{k-pk} \right) < \prod_{k=\ell+1}^j \left(1 - \frac{k-pn}{j} \right) \\ &< \exp \left(-\frac{1}{j} \sum_{k=\ell+1}^j (k-pn) \right) \leq \exp \left(-\frac{1}{2j} [(j-pn)^2 - 2] \right) \end{aligned}$$

which confirms (40). □

Lemma A.3. *Let x_1, x_2, \dots be any sequence of numbers and for each integer $j > 0$ let $y_j = \frac{1}{j} \sum_{k=1}^j x_k$. Then $y_{j+1} - y_j = \frac{1}{j+1} (x_{j+1} - y_j)$. In particular, $y_{j+1} > y_j$ if and only if $x_{j+1} > y_j$.*

Proof. The lemma follows from a simple calculation.

$$\begin{aligned}
y_{j+1} - y_j &= \left(\frac{1}{j+1} \sum_{k=1}^{j+1} x_k \right) - \left(\frac{1}{j} \sum_{k=1}^j x_k \right) \\
&= \frac{x_{j+1}}{j+1} - \left(\frac{1}{j} - \frac{1}{j+1} \right) \sum_{k=1}^j x_k \\
&= \frac{1}{j+1} \left(x_{j+1} - \frac{1}{j} \sum_{k=1}^j x_k \right) = \frac{1}{j+1} (x_{j+1} - y_j).
\end{aligned}$$

□

Proposition A.1. For a given n and p , if the value $j^* = j^*(n, p)$ defined by (6) satisfies $j^* \leq n/2$ then it satisfies

$$j^* > p(n-1) + \frac{1}{2} \sqrt{p(n-1) \ln^+ \left(\frac{p(n-1)}{16} \right)},$$

where the notation $\ln^+(x)$ is defined to be equivalent to $\max\{0, \ln(x)\}$.

Proof. We will be applying Lemma A.3 to the sequence $x_j = \binom{n-1}{j-1} p^{j-1} (1-p)^{n-j}$ and its associated sequence of averages, $y_j = \frac{1}{j} \sum_{k=1}^j x_k$. Note that $j^* = j^*(n, p)$ defined by (6) is equal to $\arg \max_j \{y_j\}$. The sequence x_k is increasing for $k < pn$, as can be seen from the calculation

$$\begin{aligned}
\frac{x_{k+1}}{x_k} &= \frac{\binom{n-1}{k} p^k (1-p)^{n-k-1}}{\binom{n-1}{k-1} p^{k-1} (1-p)^{n-k}} \\
&= \left(\frac{p}{1-p} \right) \left(\frac{(n-1)!}{k!(n-1-k)!} \right) \left(\frac{(k-1)!(n-k)!}{(n-1)!} \right) \\
&= \frac{p}{1-p} \cdot \frac{n-k}{k} = \frac{pn - pk}{k - pk}.
\end{aligned}$$

Thus, when $j < pn$, we have that x_{j+1} is greater than each of x_1, \dots, x_j hence also greater than their average, y_j . By Lemma A.3 this means that $j^* \geq pn$.

Now, by another application of Lemma A.3, the fact that $y_{j^*+1} \leq y_{j^*}$ implies that $x_{j^*+1} \leq y_{j^*}$. Noting that $y_{j^*} \leq \frac{1}{j^*}$, and recalling our assumption that $j^* \leq n/2$, we may now apply the lower bound for x_{j^*+1} supplied by Lemma A.2 to derive

$$\begin{aligned}
\frac{1}{j^*} &\geq \binom{n-1}{j^*} p^{j^*} (1-p)^{n-1-j^*} > \frac{1}{4} \cdot \sqrt{\frac{1}{j^*}} \cdot \exp \left(-\frac{2(j^* - p(n-1))^2}{p(n-1)} \right) \\
\ln(j^*) &< \ln(4) + \frac{1}{2} \ln(j^*) + \frac{2(j^* - p(n-1))^2}{p(n-1)} \\
\left[\frac{1}{2} \ln(j^*) - \ln(4) \right] p(n-1) &< 2(j^* - p(n-1))^2 \\
j^* &> p(n-1) + \sqrt{\max\{0, \left[\frac{1}{4} \ln(j^*) - \ln(2) \right]\} p(n-1)} \\
j^* &> p(n-1) + \frac{1}{2} \sqrt{\ln^+(p(n-1)/16) p(n-1)}. \tag{45}
\end{aligned}$$

□

Lemma A.4. Suppose we are given an integer n , a real number $p \in (0, 1)$, and an integer j such that $pn < j < n/2$. If $\frac{2}{j} \leq \binom{n}{j} p^j (1-p)^{n-j} \leq \frac{1}{j}$, then there exist absolute constants C_0, C_1 such that

$$\sqrt{\frac{C_0}{j \ln(j)}} \leq \sum_{k=j}^n p^k (1-p)^{n-k} \leq \sqrt{\frac{C_1}{j \ln(j)}}. \quad (46)$$

Proof. Recall from the proof of Proposition A.1 that if x_k is defined by $x_k = \binom{n}{k} p^k (1-p)^{n-k}$ then

$$\frac{x_{k+1}}{x_k} = \frac{pn - pk}{k - pk} = 1 - \frac{k - pn}{k - pk} \quad (47)$$

The identity $2x \geq -\ln(1-x)$ is valid for $0 \leq x \leq \frac{3}{4}$. So when $k - pn < \frac{3}{4}(k - pk)$ we may apply this identity to conclude

$$2 \left(\frac{k - pn}{k - pk} \right) \geq \ln x_k - \ln x_{k+1} \quad (48)$$

In particular, for $\ell = \left\lceil j + \sqrt{\frac{j}{\ln(j)}} \right\rceil$ we may sum the first inequality in (48) over $k = j, j+1, \dots, \ell-1$ to obtain

$$\ln x_j - \ln x_\ell \leq 2 \sum_{k=j}^{\ell-1} \frac{k - pn}{k - pk} \quad (49)$$

$$\leq \frac{2}{j - pj} \sum_{k=j}^{\ell-1} (k - pn) \quad (50)$$

$$= \frac{2(j - pn)(\ell - j)}{j - pj} + \frac{2}{j - pj} \sum_{k=j}^{\ell-1} (k - j) \quad (51)$$

$$< \frac{2(j - pn)(\ell - j)}{j - pj} + \frac{(\ell - j)^2}{j - pj}. \quad (52)$$

Our hypothesis that $x_j = \Theta(1/j)$ implies that $j - pn = \Theta(\sqrt{j \ln(j)})$ as was shown earlier. Hence

$$(j - pn)(\ell - j) = \Theta \left(\sqrt{j \ln(j)} \cdot \sqrt{\frac{j}{\ln(j)}} \right) = \Theta(j)$$

while $(\ell - j)^2 = \frac{j}{\ln(j)}$. Combining these bounds, and noting that the hypotheses in the lemma statement imply $p \leq 1/2$ and hence $\frac{2}{1-p} = \Theta(1)$, we find that

$$\ln x_j - \ln x_\ell = \Theta(1),$$

i.e., $x_\ell = \Theta(x_j)$. The sequence $x_j, x_{j+1}, \dots, x_\ell$ is monotonically decreasing, and we have derived that every term in this sequence is $\Theta(x_j) = \Theta(1/j)$, so we have shown that

$$x_j + \dots + x_\ell = \Theta \left(\frac{\ell - j}{j} \right) = \Theta \left(\sqrt{\frac{1}{j \ln(j)}} \right). \quad (53)$$

The binomial tail $\sum_{k>j} \binom{n}{k} p^k (1-p)^{n-k}$ is equal to $x_j + \dots + x_n$, so equation (53) implies the existence of an absolute constant C_0 such that

$$\sqrt{\frac{C_0}{j \ln(j)}} \leq \sum_{k>j} \binom{n}{k} p^k (1-p)^{n-k}, \quad (54)$$

establishing the first half of the lemma.

For the second half of the lemma, *i.e.*, the upper bound on the binomial tail, we return to (47) and apply the identity $-\ln(1-x) \geq x$, which is valid for all $x \geq 0$. Hence for $k \geq j$,

$$\ln x_k - \ln x_{k+1} \geq \frac{k-pn}{k-pk} \geq \frac{k-pn}{2(j-pj)}. \quad (55)$$

For any $\ell > j$ we may sum over $k = j, \dots, \ell - 1$ to conclude

$$\begin{aligned} \ln x_j - \ln x_\ell &\geq \frac{1}{j-pj} \sum_{k=j}^{\ell-1} (k-pn) \\ &\geq \frac{1}{j-pj} \sum_{k=j}^{\ell-1} (j-pn) \\ &= \frac{(\ell-j)(j-pn)}{j-pj} \\ \ln x_\ell &\leq \ln x_j - \left(\frac{j-pn}{j-pj} \right) (\ell-j). \end{aligned}$$

Let $s = \frac{j-pn}{j-pj}$. We have

$$\begin{aligned} \ln x_\ell &\leq \ln x_j - s(\ell-j) \\ x_\ell &\leq x_j \exp(-s(\ell-j)) \\ \sum_{\ell \geq j} x_\ell &\leq x_j \sum_{\ell=j}^{\infty} \exp(-s(\ell-j)) \\ &< x_j \left[1 + \int_0^{\infty} e^{-sx} dx \right] \\ &= x_j \left(1 + \frac{1}{s} \right). \end{aligned} \quad (56)$$

To finish, we estimate the terms appearing on the right side of (56). Recall $x_j = \Theta(1/j)$. Our estimate $j-pn = \Theta(\sqrt{j \ln(j)})$ combined with the fact that $1-p \geq \frac{1}{2}$ implies that $s = \Theta(\sqrt{\ln(j)/j})$. Therefore, $1 + \frac{1}{s} = \Theta(\sqrt{j/\ln(j)})$ and

$$\sum_{k=j}^n \binom{n}{k} p^k (1-p)^{n-k} = \sum_{\ell=j}^n x_\ell < x_j \left(1 + \frac{1}{s} \right) = \Theta \left(\frac{1}{j} \cdot \sqrt{\frac{j}{\ln(j)}} \right) = \Theta \left(\sqrt{\frac{1}{j \ln(j)}} \right),$$

which shows that the upper bound in (46) holds for a suitable choice of absolute constant C_1 . \square

Lemma A.5. For a given total prize V and participation cost c , let $\nu = V/c$ and $j = \lfloor \nu - \sqrt{\nu} \rfloor$. Consider a simple contest that splits the prize into j equal shares of size V/j . In the equilibrium of this contest, the ex-ante participation rate is greater than $\frac{1}{n-1}(\nu - \sqrt{5\nu \ln(\nu)})$.

Proof. Let $p = \frac{1}{n-1}(\nu - \sqrt{5\nu \ln(\nu)})$. In Section 3 we saw that to prove the lemma, it is sufficient to prove that

$$\begin{aligned} \frac{V}{j} \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} &\geq c \\ \sum_{k=1}^j \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} &\geq \frac{cj}{V} = \frac{j}{\nu} \leq 1 - \sqrt{\frac{1}{\nu}} \\ \sum_{k=j+1}^n \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} &\leq \sqrt{\frac{1}{\nu}}. \end{aligned} \tag{57}$$

The left side is the probability of seeing at least j heads when observing $n-1$ tosses of a coin with bias p . Letting $\lambda = p(n-1)$ denote the expected number of heads and $\delta = \frac{j}{\lambda} - 1$, the Chernoff bound (e.g. [24]) tells us that

$$\sum_{k=j+1}^n \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \leq e^{-\delta^2 \lambda / 3}.$$

To obtain an upper bound on $e^{-\delta^2 \lambda / 3}$, we derive a lower bound on $\delta^2 \lambda$ as follows.

$$\begin{aligned} \delta \lambda &= j - \lambda \geq \nu - \sqrt{\nu} - 1 - (\nu - \sqrt{5\nu \ln(\nu)}) = \sqrt{5\nu \ln(\nu)} - \sqrt{\nu} - 1 \\ \delta^2 \lambda &= \frac{(\delta \lambda)^2}{\lambda} > \frac{(\delta \lambda)^2}{\nu} \\ &= \left(\frac{\delta \lambda}{\sqrt{\nu}} \right)^2 \\ &\geq \left(\sqrt{5 \ln(\nu)} - 1 - \nu^{-1/2} \right)^2 \end{aligned}$$

The expression on the last line is clearly asymptotic to $5 \ln(\nu)$ as $\nu \rightarrow \infty$, so for all sufficiently large ν it exceeds $\frac{3}{2} \ln(\nu)$. In fact, one can verify numerically that $\left(\sqrt{5 \ln(\nu)} - 1 - \nu^{-1/2} \right)^2 > \frac{3}{2} \ln(\nu)$ for all $\nu \geq 7$. So, to prove the lemma, we divide into two cases. If $\nu < 7$ then $\nu - \sqrt{5\nu \ln(\nu)} < 0$, so the lemma's conclusion holds vacuously. If $\nu \geq 7$ then $\delta^2 \lambda > \frac{3}{2} \ln(\nu)$, so

$$\sum_{k=j+1}^n \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \leq e^{-\delta^2 \lambda / 3} < e^{-\ln(\nu)/2} = \sqrt{\frac{1}{\nu}},$$

which establishes (57) and thus completes the proof. \square

Using the foregoing lemmas, we now complete the proof of Proposition 3.3, part 3, which we restate here for convenience.

Proposition 3.3, part 3. Fix V, c , and let $j^{*,n}$ denote the number of prizes offered by the optimal contest with population size n . Let $p^{*,n}$ denote the probability that an individual agent participates in the equilibrium of that contest, and let $\lambda^{*,n} = p^{*,n} \cdot n$ denote the expected number of participants in that equilibrium. Letting $\nu = V/c$, for any $n > 2\nu$ the values of $\lambda^{*,n}$ and $j^{*,n}$ satisfy

$$\lambda^{*,n} = \nu - \Theta(\sqrt{\nu \log \nu}), \quad j^{*,n} = \nu - \Theta(\sqrt{\nu / \log \nu}). \quad (58)$$

Proof. We begin by deriving the upper and lower bounds on $\lambda = \lambda^{*,n}$. To obtain a lower bound, we use Lemma A.5, which says that a simple contest with $j = \lfloor \nu - \sqrt{\nu} \rfloor$ achieves participation rate greater than $\frac{1}{n-1}(\nu - \sqrt{5\nu \ln(\nu)})$. The optimal participation rate p can be no less than this quantity, and consequently

$$\lambda = pn > p(n-1) \geq \nu - \sqrt{5\nu \ln(\nu)}. \quad (59)$$

Before deriving an upper bound on λ we make some observations which lead to a crude lower bound on $j^* = j^{*,n}$. Observe that a winner in the optimal contest receives V/j^* and has participation cost c , so individual rationality guarantees that $V/j^* \geq c$, which we rewrite as

$$j^* \leq V/c = \nu. \quad (60)$$

By assumption, $\nu \leq n/2$, so we may apply Proposition A.1 to derive

$$j^* > \lambda + \frac{1}{2}\sqrt{\lambda \ln^+(\lambda/16)}, \quad (61)$$

where $\ln^+(x)$ is shorthand for $\max\{\ln(x), 0\}$. Combining (61) with (60) we obtain

$$\nu > \lambda + \frac{1}{2}\sqrt{\lambda \ln^+(\lambda/16)}. \quad (62)$$

When $\nu \geq 100$, the lower bound (59) implies that $\lambda \geq \nu/2$. Substituting this lower bound into (62) implies

$$\lambda < \begin{cases} \nu - \frac{1}{2}\sqrt{\frac{\nu}{2} \ln^+(\frac{\nu}{32})} & \text{if } \nu \geq 100 \\ 100 & \text{otherwise} \end{cases} \quad (63)$$

Combining (59) with (63) we obtain that $\lambda = \nu - \Theta(\sqrt{\nu \ln(\nu)})$ as $\nu \rightarrow \infty$.

As for estimating the number of prizes in the optimal contest, j^* , we begin by recalling from Proposition 3.1 that the participation rate p in the equilibrium of the optimal contest \mathcal{M} satisfies $c_{\mathcal{M}}(p) = c$. When \mathcal{M} is a simple contest that splits the prize into j^* equal shares of size j^* , we have

$$c = c_{\mathcal{M}}(p) = \frac{V}{j^*} \sum_{k=1}^{j^*} \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k}$$

and therefore

$$j^* = \frac{V}{c} \sum_{k=1}^{j^*} \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} = \nu \left[1 - \sum_{k > j^*} \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} \right]. \quad (64)$$

This shows that obtaining an asymptotic estimate of $\nu - j^*$ as a function of ν (as $\nu \rightarrow \infty$) is equivalent to estimating the binomial tail $\sum_{k>j^*} \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k}$ in terms of ν . Lemma A.4 establishes that

$$\sum_{k>j^*} \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k} = \Theta \left(\sqrt{\frac{1}{j^* \ln(j^*)}} \right) \quad (65)$$

and substituting this estimate in (64) yields

$$j^* = \nu - \Theta \left(\sqrt{\frac{\nu}{j^* \ln(j^*)}} \right). \quad (66)$$

We have already shown above that for sufficiently large ν , $\lambda = \Theta(\nu)$ and $j^* = \Theta(\lambda) = \Theta(\nu)$. Thus, $j^* \ln(j^*) = \Theta(\nu \ln(\nu))$, and we can substitute this estimate into (66) to derive $j^* = \nu - \Theta(\sqrt{\nu / \ln(\nu)})$. \square

B Equilibrium Existence and Uniqueness

In this section we restate and prove three lemmas related to existence and uniqueness of symmetric equilibria in rank-order contests. Throughout this section we assume that the distribution $F(q, c)$ is absolutely continuous with respect to Lebesgue measure.

Lemma B.1. *Any rank-order contest \mathcal{M} has a symmetric mixed-strategy equilibrium.*

Proof. The lemma essentially follows from Milgrom and Weber’s 1985 mixed equilibrium existence theorem for games of incomplete information. Their theorem requires two hypotheses: equicontinuous payoffs (R1) and absolutely continuous information (R2). The game induced by contest \mathcal{M} is easily seen to satisfy both of these hypotheses. R1 is satisfied because each player has only two actions (Proposition 1(a) of [20]) and R2 is satisfied because agents’ types are independent (Proposition 3(b) of [20]). The conclusion of their theorem asserts that a mixed-strategy equilibrium exists and does not specify existence of a symmetric equilibrium for symmetric games such as ours. However, the existence of a symmetric equilibrium is established by applying Glicksberg’s Fixed Point Theorem [12] in exactly the same way that it is applied in the proof of Theorem 1 of [20], except that the fixed point theorem is applied to the point-to-convex-set mapping that maps every mixed strategy μ (of a single agent) to that agent’s best response set when responding to the symmetric strategy profile in which all agents play x . \square

Lemma B.2. *If μ is any symmetric mixed strategy equilibrium of a rank-order contest \mathcal{M} then there exists a symmetric pure strategy equilibrium π that is “almost equal” to μ in the following sense: there exists a measure-zero set of types, T , such that for all $(q, c) \notin T$, when π' is randomly sampled from μ , the event $\pi'(q, c) \neq \pi(q, c)$ has probability zero.*

Proof. Define F_μ to be the cumulative distribution function of the random variable $q \cdot \pi'(q, c)$ when (q, c) is drawn from F and π' is drawn from μ . As in the proof of Lemma 4.1, an agent of type (q, c) facing $n - 1$ opponents all of whom are playing mixed strategy μ receives an expected prize equal to $c_{\mathcal{M}}(1 - F_\mu(q))$ if she participates. If $c > c_{\mathcal{M}}(1 - F_\mu(q))$ then the agent strictly prefers non-participation to participation, hence $\pi'(q, c) = 0$ for almost every π' in μ . If $c < c_{\mathcal{M}}(1 - F_\mu(q))$ then the agent strictly prefers participation to non-participation, hence $\pi'(q, c) = 1$ for almost every π' in μ . This verifies that the lemma’s conclusion holds with

$$\pi(q, c) = \begin{cases} 1 & \text{if } c < c_{\mathcal{M}}(1 - F_\mu(q)) \\ 0 & \text{if } c \geq c_{\mathcal{M}}(1 - F_\mu(q)) \end{cases}$$

and $T = \{(q, c_{\mathcal{M}}(1 - F_{\mu}(q)))\}$. This set T has measure zero because it intersects each vertical line $\{q\} \times \mathbb{R}$ in a single point. \square

Lemma B.3. *If \mathcal{M} is any rank-order contest and π, π' are two of its symmetric pure strategy equilibria then $\{(q, c) \mid \pi(q, c) \neq \pi'(q, c)\}$ has measure zero.*

Proof. Since π, π' are both equilibria, they are also both sub-equilibria. Let $F_{\pi}, F_{\pi'}$ denote the cumulative distribution functions of $q \cdot \pi(q, c)$ and $q \cdot \pi'(q, c)$, respectively. Lemma 4.1 applied to the sub-equilibrium π implies $F_{\pi}(x) \geq F_{\pi'}(x)$ for all x . Applying the same lemma to the sub-equilibrium π' yields the reverse inequality. Hence $F_{\pi} = F_{\pi'}$. Reasoning now as in the proof of Lemma ??, we may conclude that $\pi(q, c) = \pi'(q, c)$ unless $c = c_{\mathcal{M}}(1 - F_{\pi}(q))$. \square

C Proofs from Section 4.2

This section contains the proofs comparing the expectation of the maximum quality elicited by the winner-take-all mechanism to that of the optimal mechanism. Recall that the overall structure of the proof divides the participants in the optimal mechanism into two subpopulations—those with cost greater than $\frac{V}{2}$ and those with cost less than or equal to $\frac{V}{2}$. The following lemmas treat those sub-populations separately.

Lemma C.1. $\text{WTA} \geq \frac{1}{2}E[\max\{q_i \mid i \in \mathcal{E}^*; c_i \leq \frac{V}{2}\}]$.

Proof. First, note that $E[\max\{q_i \mid i \in \mathcal{E}^*; c_i \leq \frac{V}{2}\}] \leq E[\max\{q_i \mid c_i \leq \frac{V}{2}\}]$, so it is sufficient to show that $\text{WTA} \geq E[\max\{q_i \mid c_i \leq \frac{V}{2}\}]$. This allows us to reason only about strategic behavior in the winner-take-all mechanism, instead of needing to also understand participation choices in the optimal mechanism. We will use the sub-equilibrium lemma to relate equilibrium outcomes in \mathcal{M}_{WTA} to the expected output in a particular sub-equilibrium, identified below, that can be compared against $E[\max\{q_i \mid c_i \leq \frac{V}{2}\}]$.

Consider $n - 1$ random draws of (q_i, c_i) from F , and let μ denote the *median* of the random variable $\max\{q_i \mid c_i \leq \frac{V}{2}\}$. That is, μ is such that:

$$\Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > \mu) = \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} \leq \mu). \quad (67)$$

Consider the strategy profile π_1^S where an agent i participates in \mathcal{M}_{WTA} if and only if $q_i \geq \mu$ and $c_i \leq \frac{V}{2}$. We claim that π_1^S constitutes a sub-equilibrium of \mathcal{M}_{WTA} : by definition of μ , if i participates when playing π_1^S , her expected payoff in \mathcal{M}_{WTA} is

$$E[u_i] = V \cdot \Pr(q_i \geq \max_{j \neq i} q_j \mid j \in \mathcal{S}(\pi_1^S)) - c \geq V \cdot \frac{1}{2} - \frac{V}{2} \geq 0,$$

so that it is individually rational for each agent to participate if all other agents play according to π_1^S .

We now analyze the expected outcome in this sub-equilibrium π_1^S of \mathcal{M}_{WTA} :

$$\begin{aligned}
E[\max_n \{q_i \mid q_i \geq \mu; c_i \leq \frac{V}{2}\}] &= \int_0^\infty \Pr(\max\{q_i \mid q_i \geq \mu; c_i \leq \frac{V}{2}\} > t) dt \\
&= \int_0^\mu \Pr(\exists q_i \mid q_i \geq \mu; c_i \leq \frac{V}{2}) dt + \int_\mu^\infty \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > t) dt \\
&= \mu \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} \geq \mu) + \int_\mu^\infty \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > t) dt \\
&\geq \mu \cdot \frac{1}{2} + \frac{1}{2} \int_\mu^\infty \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > t) dt \\
&= \frac{1}{2} \left(\int_0^\mu 1 \cdot dt + \int_\mu^\infty \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > t) dt \right) \\
&\geq \frac{1}{2} \left(\int_0^\mu \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > t) dt + \int_\mu^\infty \Pr(\max\{q_i \mid c_i \leq \frac{V}{2}\} > t) dt \right) \\
&= \frac{1}{2} E \max\{q_i \mid c_i \leq \frac{V}{2}\},
\end{aligned}$$

where the inequality in the fourth line follows from the definition of μ in (67), and because μ depends on the max of $n - 1$ random draws of q_i while the max here is over n draws from q_i .

By the corollary to the sub-equilibrium lemma (Corollary 4.1), the expected maximum output in any equilibrium of \mathcal{M}_{WTA} is at least as large as the expected maximum output in a sub-equilibrium, and therefore specifically in the sub-equilibrium π_1^S . Combining this with the inequality just derived, we have

$$\text{WTA} \geq E[\max_n \{q_i \mid q_i \geq \mu; c_i \leq \frac{V}{2}\}] \geq \frac{1}{2} E[\max_n \{q_i \mid c_i \leq \frac{V}{2}\}] \geq \frac{1}{2} E[\max\{q_i \mid i \in \mathcal{E}^*; c_i \leq \frac{V}{2}\}],$$

which proves the lemma. \square

Lemma C.2. $\text{WTA} \geq E[\max\{q_i \mid i \in \mathcal{E}^*; c_i > \frac{V}{2}\}].$

Proof. To prove this inequality, we use the sub-equilibrium lemma (Lemma 4.1) again, albeit with a different strategy profile. Consider the equilibrium $\pi^E(\mathcal{M}^*)$ of an optimal mechanism \mathcal{M}^* , and let \mathcal{E}^* denote the set of participants in \mathcal{M}^* when agents play according to $\pi^E(\mathcal{M}^*)$ as before. Consider a strategy π_2^S for \mathcal{M}_{WTA} where an agent participates if and only if she has cost $c > \frac{V}{2}$, and would also participate according to $\pi^E(\mathcal{M}^*)$. We claim that π_2^S is a sub-equilibrium of the winner-take-all mechanism \mathcal{M}_{WTA} , i.e., that the strategy profile where all agents in $\mathcal{S}(\pi_2^S) = \{i \mid i \in \mathcal{E}^*; c_i > \frac{V}{2}\}$ participate and all remaining agents do not participate leads to non-negative payoffs for all agents in the *winner-take-all* mechanism.

To show this, consider the following *lottery over simple* mechanisms defined by the rewards in the mechanism $\mathcal{M}^*(v_1, \dots, v_n)$: the lottery runs the simple contest which awards equal prizes of $\frac{V}{j}$ each to the top j contestants with probability

$$\Pr(j) = \frac{j}{V}(v_j - v_{j+1}).$$

Note that (i) $\sum_{j=1}^n \Pr(j) = 1$ since $\sum_{j=1}^n v_j = V$ (this is derived in Equation (4) in §3 for *any* feasible mechanism) so that the distribution $\Pr(j)$ is a valid probability distribution, and (ii) for any draw of the n values of (q_i, c_i) , the expected payoff to the agent with the j^{th} highest quality is exactly equal to v_j in this lottery over simple mechanisms. That is, this lottery over simple mechanisms is payoff-equivalent (in expectation) for all agents to the optimal mechanism $\mathcal{M}^*(v_1, \dots, v_n)$.

The expected payoff of an agent in \mathcal{E}^* (i.e., a participating agent) in the *optimal* mechanism $\mathcal{M}^*(v_1, \dots, v_n)$ can therefore be written as

$$E[u_i] = \sum_{j=1}^n \frac{j(v_j - v_{j+1})}{V} \left(\Pr(i \text{ belongs to the top } j \text{ agents in } \mathcal{E}^*) \cdot \frac{V}{j} - c_i \right).$$

Since π^E is an equilibrium of \mathcal{M}^* , every agent who participates must receive non-negative expected payoff. Specifically, consider the agents in \mathcal{E}^* with cost greater than $V/2$. All but the term corresponding to $j = 1$ in the summation are negative for such an agent, since $\frac{V}{j} \leq \frac{V}{2} < c_i$ for $j \geq 2$. So the term corresponding to $j = 1$ must be positive. But for $j = 1$, the expression $\left[\Pr(i \text{ belongs to the top } j \text{ agents in } \mathcal{E}^*) \cdot \frac{V}{j} - c_i \right]$ is exactly the expected payoff of an agent in \mathcal{E}^* in the *winner-take-all* mechanism \mathcal{M}_{WTA} , if all agents in \mathcal{E}^* participate and all agents not in \mathcal{E}^* do not. Since the expected payoff to this agent when only agents in $\mathcal{E}^* \cap \{i \mid c_i > \frac{V}{2}\}$ participate can only be larger than this value (she is more likely to be the highest-quality agent when competing with a smaller subset of agents), this means that the strategy profile π_2^S where all agents in \mathcal{E}^* with $c_i > \frac{V}{2}$ participate, and the remaining agents do not, is a sub-equilibrium of the winner-take-all mechanism.

Applying Corollary 4.1 to this strategy profile π_2^S , we have $\text{WTA} \geq E[\max\{q_i \mid i \in \mathcal{E}^*; c_i > \frac{V}{2}\}]$, which proves the lemma. \square